Revisiting Decomposable Submodular Function Minimization with Incidence Relations

Pan Li
UIUC
panli2@illinois.edu

Olgica Milenkovic
UIUC
milenkov@illinois.edu

Abstract

We introduce a new approach to decomposable submodular function minimization (DSFM) that exploits incidence relations. Incidence relations describe which variables effectively influence the component functions, and when properly utilized, they allow for improving the convergence rates of DSFM solvers. Our main results include the precise parametrization of the DSFM problem based on incidence relations, the development of new scalable alternative projections and parallel coordinate descent methods and an accompanying rigorous analysis of their convergence rates.

1 Introduction

A set function $F : 2^{[N]} \to \mathbb{R}$ over a ground set $[N]$ is termed submodular if for all pairs of sets $S_1, S_2 \subseteq [N]$, one has $F(S_1) + F(S_2) \geq F(S_1 \cap S_2) + F(S_1 \cup S_2)$. Submodular functions capture the ubiquitous phenomenon of diminishing marginal costs [1] and they frequently arise as part of the objective function of various machine learning optimization problems [2, 3, 4, 5, 6, 7].

Among the various submodular function optimization problems, submodular function minimization (SFM), which may be stated as $\min_{S \subseteq [N]} F(S)$, is one of the most important and commonly studied questions. The current fastest known SFM algorithm has complexity $O(N^4 \log^{O(1)} N + \tau N^3)$, where τ denotes the time needed to evaluate the submodular function [8]. Although SFM solvers operate in time polynomial in N, the high-degree of the underlying polynomial prohibits their use in practical large-scale settings. For this reason, a recent line of work has focused on developing scalable and parallelizable algorithms for solving the SFM problem by leveraging the property of *decomposability* [9]. Decomposability asserts that the submodular function may be written as a sum of "simpler" submodular functions that may be optimized sequentially or in parallel. Formally, the underlying problem, referred to as decomposable SFM (DSFM), may be stated as:

$$\text{DSFM:} \quad \min_{S} \sum_{r \in [R]} F_r(S), \tag{1}$$

where $F_r : 2^{[N]} \to \mathbb{R}$ is a submodular function for all $r \in [R]$. Algorithmic solutions for the DSFM problem fall into two categories, combinatorial optimization approaches [10, 11] and continuous function optimization methods [12]. In the latter setting, a crucial concept is the Lovász extension of the submodular function which is convex [13] and lends itself to a norm-regularized convex optimization framework. Prior work in continuous DSFM has focused on devising efficient algorithms for solving the convex problem and deriving matching convergence results. The best known approaches include the alternating projection (AP) methods [14, 15] and the coordinate descent (CD) methods [16].

Despite some simplifications offered through decomposibility, DSFM algorithms still suffer from scalability issues and have convergence guarantees that are suboptimal. To address the first issue, one needs to identify additional problem constraints that allow for parallel implementations. To resolve the second issue and more precisely characterize and improve the convergence rates, one needs to better understand how the individual submodular components jointly govern the global optimal solution. In both cases, it is crucial to utilize *incidence relations* that describe which subsets of variables

32nd Conference on Neural Information Processing Systems (NeurIPS 2018), Montréal, Canada.

directly affect the value of any given component function. Often, incidences involve relatively small subsets of elements, which leads to desirable sparsity constraints. This is especially the case for min-cut problems on graphs and hypergraphs (where each submodular component involves two or several vertices) [17, 18] and MAP inference with higher-order potentials (where each submodular component involves variables corresponding to adjacent pixels) [9]. Although incidence relations have been used to parametrize the algorithmic complexity of combinatorial optimization methods for solving DSFM problems [10], they have been largely overlooked in continuous optimization methods. Some prior work considered merging decomposable parts with nonoverlapping support into one submodular function, thereby creating a coarser decomposition that may be processed more efficiently [14, 15, 16], but the accompanying algorithms were neither designed in a form that can optimally use this information nor analyzed precisely with respect to their convergence rates and merging strategies. In an independent work, Djolonga and Krause found that the variational inference problem in L-FIELD could be reduced to a DSFM problem with sparse incidence relations [19], while their analysis only worked for regular cases.

Here, we revisit two benchmark algorithms for continuous DSFM – AP and CD – and describe how to modify them to exploit incidence relations that allow for significantly improved computational complexity. Furthermore, we provide a complete theoretical analysis of the algorithms parametrized by incidence relations with respect to their convergence rates. AP-based methods that leverage incidence relations achieve better convergence rates than classical AP algorithms both in the sequential and parallel optimization scenario. The random CD method (RCDM) and accelerated CD method (ACDM) that incorporate incidence information can be parallelized. The complexity of sequential CD methods cannot be improved using incidence relations, but the convergence rate of parallel CD methods strongly depends on how the incidence relations are used for coordinate sampling: while a new specialized combinatorial sampling based on equitable coloring [20] is optimal, uniformly at random sampling produces a 2-approximation. It also leads to a greedy method that empirically outperforms random sampling. A summary of these and other findings is presented in Table 1.

	Prior work		This work	
	Sequential	Parallel	Sequential	Parallel
AP	$O(N^2R^2)$	$O(N^2R^2/K)$	$O(N\|\mu\|_1 R)$	$O(N\|\mu\|_1 R/K)$
RCDM	$O(N^2R)$	-	$O(N^2R)$	$O\left(\left(\frac{R-K}{R-1}N^2 + \frac{K-1}{R-1}N\|\mu\|_1\right)R/K\right)$
ACDM	$O(NR)$	-	$O(NR)$	$O\left(\left(\frac{R-K}{R-1}N^2 + \frac{K-1}{R-1}N\|\mu\|_1\right)^{1/2}R/K\right)$

Table 1: Overview of known and new results: each entry contains the required number of iterations to achieve an ϵ-optimal solution (the dependence on ϵ is the same for all algorithms and hence omitted). Here, $\|\mu\|_1 = \sum_{i\in[N]} \mu_i$, where for all $i \in [N]$, μ_i equals the number of submodular functions that involve element i; K is a parallelization parameter that equals the number of min-norm points problems that have to be solved within each iteration.

2 Background, Notation and Problem Formulation

We start our exposition by reviewing several recent lines of work for solving the DSFM problem, and focus on approaches that transform the DSFM problem into a continuous optimization problem. Such approaches exploit the fact that the Lovász extension of a submodular function is convex. Without loss of generality, we tacitly assume that all submodular functions F_r are normalized, i.e., that $F_r(\emptyset) = 0$ for all $r \in [R]$. Also, we define given a vector $z \in \mathbb{R}^N$ and $S \subseteq [N]$, $z(S) = \sum_{i\in S} z_i$. Then, the *base polytope* of the r-th submodular function F_r is defined as

$$\mathcal{B}_r \triangleq \{y_r \in \mathbb{R}^N | y_r(S) \leq F_r(S), \text{ for any } S \subset [N], \text{and } y_r([N]) = F_r([N])\}.$$

The *Lovász extention* [13] $f_r(\cdot) : \mathbb{R}^N \to \mathbb{R}$ of a submodular function F_r is defined as $f_r(x) = \max_{y_r\in\mathcal{B}_r} \langle y_r, x\rangle$, where $\langle\cdot,\cdot\rangle$ denotes the inner product of two vectors. The DSFM problem can be solved through continuous optimization, $\min_{x\in[0,1]^N} \sum_r f_r(x)$. To counter the nonsmoothness of the objective function, a proximal formulation of a generalization of the above optimization problem is considered instead [14],

$$\min_{x\in\mathbb{R}^N} \sum_{r\in[R]} f_r(x) + \frac{1}{2}\|x\|_2^2. \tag{2}$$

As the problem (2) is strongly convex, it has a unique optimal solution, denoted by x^*. The exact discrete solution to the DSFM problem equals $S^* = \{i \in [N] | \, x_i^* > 0\}$.

For convenience, we denote the product of base polytopes as $\mathcal{B} = \otimes_{r=1}^{R} \mathcal{B}_r$, and write $y = (y_1, y_2, ..., y_R) \in \mathcal{B}$. Also, we let A be a simple linear mapping $\otimes_{r=1}^{R} \mathbb{R}^N \to \mathbb{R}^N$, which given a point $a = (a_1, a_2, ..., a_R) \in \otimes_{r=1}^{R} \mathbb{R}^N$ outputs $Aa = \sum_{r \in [R]} a_r$. The AP and CD algorithms for solving (2) use the dual form of the problem, described in the next lemma.

Lemma 2.1 ([14])**.** The dual problem of (2) reads as

$$\min_{a,y} \|a - y\|_2^2 \quad \text{s.t.} \quad Aa = 0, \; y \in \mathcal{B}. \tag{3}$$

Moreover, problem (3) may be written in the more compact form

$$\min_{y} \|Ay\|_2^2 \quad \text{s.t.} \quad y \in \mathcal{B}. \tag{4}$$

For both problems, the primal and dual variables are related according to $x = -Ay$. In what follows, for notational simplicity, we write $g(y) = \frac{1}{2}\|Ay\|_2^2$.

The AP [15] and RCD algorithms [16] described below provide solutions to the problems (3) and (4), respectively. They both rely on repeated projections $\Pi_{\mathcal{B}_r}(\cdot)$ onto the base polytopes $\mathcal{B}_r$, $r \in [R]$. These projections are typically less computationally intense than projections onto the complete base polytope of F as they involve fewer data dimensions. The projection operation $\Pi_{\mathcal{B}_r}(\cdot)$ requires one to solve a min-norm problem by either exploiting the special forms of F_r or by using the general purpose algorithm of Wolfe [21]. The complexity of the method is typically characterized by the number of required projections $\Pi_{\mathcal{B}_r}(\cdot)$.

The AP algorithm. Starting with $y = y^{(0)}$, iteratively compute a sequence $(a^{(k)}, y^{(k)})_{k=1,2,...}$ such that for all $r \in [R]$, $a_r^{(k)} = y_r^{(k-1)} - Ay^{(k-1)}/R$, $y_r^{(k)} = \Pi_{\mathcal{B}_r}(a_r^{(k)})$, until a stopping criteria is met.

The RCDM algorithm. In each iteration k, chose uniformly at random a subset of elements in y associated with one atomic function in the decomposition (1), say the one with index r_k. Then, compute the sequence $(y^{(k)})_{k=1,2,...}$ according to $y_{r_k}^{(k)} = \Pi_{B_{r_k}}\left(-\sum_{r \neq r_k} y_r^{(k-1)}\right)$, $y_r^{(k)} = y_r^{(k-1)}$, for $r \neq r_k$.

Finding an ϵ-optimal solution for both the AP and RCD methods requires $O(N^2 R \log(\frac{1}{\epsilon}))$ iterations. In each iteration, the AP algorithm computes the projections onto all R base polytopes, while the RCDM only computes one projection. Therefore, as may be seen from Table 1, the sequential AP solver, which computes one projection in each iteration, requires $O(N^2 R^2 \log(\frac{1}{\epsilon}))$ iterations. However, the projections within one iteration of the AP method can be generated in parallel, while the projections performed in the RCDM have to be generated sequentially.

2.1 Incidence Relations and Related Notations

We next formally introduce one of the key concepts used in this work: *incidence relations* between elements of the ground set and the component submodular functions.

We say that an element $i \in [N]$ is *incident* to a submodular function F iff there exists a $S \subseteq [N]/\{i\}$ such that $F(S \cup \{i\}) \neq F(S)$; similarly, we say that the submodular function F is *incident* to an element i iff i is incident to F. To verify whether an element i is incident to a submodular function F, one needs to verify that $F(\{i\}) = 0$ and that $F([N]) = F([N]/\{i\})$ since for any $S \subseteq [N]/\{i\}$

$$F(\{i\}) \geq F(S \cup \{i\}) - F(S) \geq F([N]) - F([N]/\{i\}).$$

Furthermore, note that if $i \in [N]$ is not incident to F_r, then for any $y_r \in \mathcal{B}_r$, one has $y_{r,i} = 0$. Let S_r be the set of all elements incident to F_r. For each element i, denote the number of submodular functions that are incident to i by $\mu_i = |\{r \in [R] : i \in S_r\}|$. We also refer to μ_i as the degree of element i. We find it useful to partition the set of submodular functions into different groups. Given a group $C \subseteq [R]$ of submodular functions, we define the degree of the element i within C, μ_i^C, as $\mu_i^C = |\{r \in C : i \in S_r\}|$.

We also define a skewed norm involving two vectors $w \in \mathbb{R}_{>0}^N$ and $z \in \mathbb{R}^N$ according to $\|z\|_{2,w} \triangleq \sqrt{\sum_{i \in [N]} w_i z_i^2}$. With a slight abuse of notation, for two vectors $\theta = (\theta_1, \theta_2, ..., \theta_R) \in \otimes_{r=1}^{R} \mathbb{R}_{>0}^N$

and $y \in \otimes_{r=1}^{R}\mathbb{R}^N$, we also define the norm $\|y\|_{2,\theta} \triangleq \sqrt{\sum_{r\in[R]} \|y_r\|_{2,\theta_r}^2}$. Which of the norms we refer to should be clear from the context. In addition, we let $\|\theta\|_{1,\infty} = \sum_{i\in[N]} \max_{r\in[R]:i\in S_r} \theta_{r,i}$. For a closed set $\mathcal{K} \subseteq \otimes_{r=1}^{R}\mathbb{R}^N$ and a positive vector $\theta \in \otimes_{r=1}^{R}\mathbb{R}^N_{>0}$, the distance between y and $\mathcal{K}$ is defined as $d_\theta(y,\mathcal{K}) = \min\{\|y-z\|_{2,\theta} | z \in \mathcal{K}\}$. Also, given a set $\Omega \subseteq \mathbb{R}^N$, we let $\Pi_{\Omega,w}(\cdot)$ denote the projection operation onto Ω with respect to the norm $\|\cdot\|_{2,w}$.

Given a vector $w \in \mathbb{R}^N_{>0}$, we also make use of an induced vector $I(w) \in \otimes_{r=1}^{R}\mathbb{R}^N$ whose r-th entry satisfies $(I(w))_r = w$. It is easy to check that $\|I(w)\|_{1,\infty} = \|w\|_1$. Of special interest are induced vectors based on pairs of N-dimensional vectors, $\mu = (\mu_1, \mu_2, ..., \mu_N)$, $\mu^C = (\mu_1^C, \mu_2^C, ..., \mu_N^C)$. Finally, for $w, w' \in \mathbb{R}^N$, we denote the element-wise power of w by $w^\alpha = (w_1^\alpha, w_2^\alpha, ..., w_N^\alpha)$, for some $\alpha \in \mathbb{R}$, and the element-wise product of w and w' by $w \odot w' = (w_1 w_1', w_2 w_2', ..., w_N w_N')$.

Next, recall that x^* is the unique optimal solution of the problem (2) and let $\mathcal{Z} = \{\xi \in \otimes_{r=1}^{R}\mathbb{R}^N | A\xi = -x^*, \xi_{r,i} = 0, \forall i \in S_r, \forall r \in [R]\}$. Then, due to the duality relationship of Lemma 2.1, $\Xi = \mathcal{Z} \cap \mathcal{B}$ is the set of optimal solutions $\{y\}$.

3 Continuous DSFM Algorithms with Incidence Relations

In what follows, we revisit the AP and CD algorithms and describe how to improve their performance and analytically establish their convergence rates. Our first result introduces a modification of the AP algorithm (3) that exploits incidence relations so as to decrease the required number of iterations from $O(N^2R)$ to $O(N\|\mu\|_1)$. Our second result is an example that shows that the convergence rates of CD algorithms [11] cannot be directly improved by exploiting the functions' incidence relations even when the incidence matrix is extremely sparse. Our third result is a new algorithm that relies of coordinate descent steps but can be parallelized. In this setting, incidence relations are essential to the parallelization process.

To analyze solvers for the continuous optimization problem (2) that exploit the incidence structure of the functions, we make use of the skewed norm $\|\cdot\|_{2,w}$ with respect to some positve vector w that accounts for the fact that incidences are, in general, nonuniformly distributed. In this context, the projection $\Pi_{\mathcal{B}_r,w}(\cdot)$ reduces to solving a classical min-norm problem after a simple transformation of the underlying space which does not incur significant complexity overheads. To see this, note that in order to solve a generic min-norm point problem, one typically uses either Wolfe's algorithm (continuous) or a divide-and-conquer procedure (combinatorial). The complexity of the former is at most quadratic in $F_{r,\max} \triangleq \max_{v,S} |F_r(S\cup\{v\}) - F_r(S)|$ [22], while the complexity of the latter merely depends on $\log F_{r,\max}$ [14] (see Section A in the Supplement). It is unclear if including the weight vector w into the projection procedure increases or decreases $F_{r,\max}$. In either case, given that in our derivations all elements of w are contained in $[1, \max_{i\in[N]} \mu_i]$ instead of N or R, we do not expect to see significant changes in the complexity of the projection operation. Hence, throughout the remainder of our exposition, we regard the projection operation as an oracle and measure the complexity of all algorithms in terms of the number of projections performed.

Also, observe that one may avoid computing projections in skewed-norm spaces by introducing in (2) a weighted rather than an unweighted proximal term. This gives another continuous objective that still provides a solution to the discrete problem (1). Even in this case, we can prove that the numbers of iterations used in the different methods listed Table 1 remain the same. Furthermore, by combining projections in skewed-norm spaces and weighted proximal terms, it is possible to actually reduce the number of iterations given in Table 1. However, for simplicity, we focus on the objective (2) and projections in skewed-norm spaces. Methods using weighted proximal terms with and without skewed-norm projections are analyzed in a similar manner in Section L of the Supplement.

We make frequent use of the following result which generalizes Lemma 4.1 of [11].

Lemma 3.1. Let $\theta \in \otimes_{r=1}^{R}\mathbb{R}^N_{>0}, w \in \mathbb{R}^N_{>0}$ be two positive vectors. Let $y \in \mathcal{B}$ and let z be in the base polytope of the submodular function F. Then, there exists a point $\xi \in \mathcal{B}$ such that $A\xi = z$ and $\|\xi - y\|_{2,\theta} \leq \sqrt{\frac{\|\theta\|_{1,\infty}}{2}} \|Ay - z\|_1$. Moreover, $\|\xi - y\|_{2,\theta} \leq \sqrt{\frac{\|\theta\|_{1,\infty}\|w^{-1}\|_1}{2}} \|Ay - z\|_{2,w}$.

3.1 The Incidence Relation AP (IAP)

The following result establishes the basis of our improved AP method leveraging incidence structures.

Lemma 3.2. The following problem is equivalent to problem (3):

$$\min_{a,y} \|a-y\|^2_{2,I(\mu)} \quad \text{s.t.} \quad y \in \mathcal{B}, Aa = 0, \text{ and } a_{r,i} = 0, \ \forall (r,i) : i \notin S_r, r \in [R]. \tag{5}$$

Let $\mathcal{A} = \{a \in \otimes_{r=1}^{R} \mathbb{R}^N | Aa = 0, a_{r,i} = 0, \ \forall (r,i) : i \notin S_r\}$ and $\mathcal{A}' = \{a \in \otimes_{r=1}^{R} \mathbb{R}^N | Aa = 0\}$. The AP algorithm for problem (5) consists of alternatively computing projections between $\mathcal{A}$ and $\mathcal{B}$, as opposed to those between $\mathcal{A}'$ and $\mathcal{B}$ used in the problem (3). However, as already pointed out, unlike for the classical AP problem (3), the distance in (5) is not Euclidean, and hence the projections may not be orthogonal.

The IAP method for solving (5) proceeds as follows. We begin with $a = a^{(0)} \in \mathcal{A}$, and iteratively compute a sequence $(a^{(k)}, y^{(k)})_{k=1,2,\ldots}$ as follows: for all $r \in [R]$, $y_r^{(k)} = \Pi_{\mathcal{B}_r,\mu}(a_r^{(k)})$, $a_{r,i}^{(k)} = y_{r,i}^{(k-1)} - \mu_i^{-1}(Ay^{(k-1)})_i$, $\forall\, i \in S_r$. The key difference between the AP and IAP algorithms is that the latter effectively removes "irrelevant" components of y_r by fixing the irrelevant components of a to 0. In the AP method of Nishihara [15], these components are never zero as they may be "corrupted" by other components during AP iterations. Removing irrelevant components results in projecting y into a subspace of lower dimensions, which significantly accelerates the convergence of IAP.

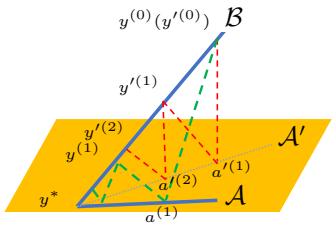

Figure 1: Illustration of the IAP method for solving problem (5): The space $\mathcal{A}$ is a subspace of $\mathcal{A}'$, which leads to faster convergence of the IAP method when compared to AP.

The analysis of the convergence rate of the IAP method follows a similar outline as that used to analyze (3) in [15]. Following Nishihara et al. [15], we define the following parameter that plays a key role in determining the rate of convergence of the AP algorithm, $\kappa_* \triangleq \sup_{y \in \mathcal{Z} \cup \mathcal{B}/\Xi} \frac{d_{I(\mu)}(y,\Xi)}{\max\{d_{I(\mu)}(y,\mathcal{Z}), d_{I(\mu)}(y,\mathcal{B})\}}$.

Lemma 3.3 ([15])**.** If $\kappa_* < \infty$, the AP algorithm converges linearly with rate $1 - \frac{1}{\kappa_*^2}$. At the k-th iteration, the algorithm outputs a value $y^{(k)}$ that satisfies

$$d_{I(\mu)}(y^{(k)}, \Xi) \leq 2 d_{I(\mu)}(y^{(0)}, \Xi) \left(1 - \frac{1}{\kappa_*^2}\right)^k.$$

To apply the above lemma in the IAP setting, one first needs to establish an upper bound on κ_*. This bound is given in Lemma 3.4 below.

Lemma 3.4. The parameter κ_* is upper bounded as $\kappa_* \leq \sqrt{N\|\mu\|_1/2} + 1$.

By using the above lemma and the bound on κ_*, one can establish the following convergence rate for the IAP method.

Theorem 3.5. After $O(N\|\mu\|_1 \log(1/\epsilon))$ iterations, the IAP algorithm for solving problem (5) outputs a pair of points (a, y) that satisfies $d_{I(\mu)}(y, \Xi) \leq \epsilon$.

Note that in practice, one often has $\|\mu\|_1 \ll NR$, which shows that the convergence rate of the AP method for solving the DSBM problem may be significantly improved.

3.2 Sequential Coordinate Descent Algorithms

Unlike the AP algorithm, the CD algorithms by Ene et al. [16] remain unchanged given (4). Our first goal is to establish whether the convergence rate of the CD algorithms can be improved using a parameterization that exploits incidence relations.

The convergence rate of CD algorithms is linear if the objective function is component-wise smooth and ℓ-strong convex. In our case, $g(y)$ is component-wise smooth as for any $y, z \in \mathcal{B}$ that only differ

in the r-th block (i.e., $y_r \neq z_r$, $y_{r'} = z_{r'}$ for $r' \neq r$), one has

$$\|\nabla_r g(y) - \nabla_r g(z)\|_2 \leq \|y - z\|_2. \tag{6}$$

Here, $\nabla_r g$ denotes the gradient vector associated with the r-th block.

Definition 3.6. We say that the function $g(y)$ is *ℓ-strongly convex* in $\|\cdot\|_2$, if for any $y \in \mathcal{B}$

$$g(y^*) \geq g(y) + \langle \nabla g(y), y^* - y \rangle + \frac{\ell}{2}\|y^* - y\|_2^2, \text{ or equivalently, } \quad \|Ay - Ay^*\|_2^2 \geq \ell \|y^* - y\|_2^2,$$

where $y^* = \arg\min_{z \in \Xi} \|z - y\|_2^2$. Moreover, we let $\ell_* = \sup\{\ell : g(y) \text{ is } \ell\text{-strongly convex in } \|\cdot\|_2\}$.

Note that the above definition essentially establishes a form of weak-strong convexity [23]. Then, using standard analytical tools for CD algorithms [24], we can prove the following result [16].

Theorem 3.7. The RCDM for problem (4) outputs a point y that satisfies $\mathbb{E}[g(y)] \leq g(y^*) + \epsilon$ after $O(\frac{R}{\ell_*}\log(1/\epsilon))$ iterations. The ACDM applied to the problem (4) outputs a point y that satisfies $\mathbb{E}[g(y)] \leq g(y^*) + \epsilon$ after $O(\frac{R}{\sqrt{\ell_*}}\log(1/\epsilon))$ iterations.

To precisely characterize the convergence rate, we need to find an accurate estimate of ℓ_*. Ene et al. [11] derived $\ell_* \geq \frac{1}{N^2}$ without taking into account the incidence structure. As sparse incidence side information improves the performance of the AP method, it is of interest to determine if the same can be accomplished for the CD algorithms. Example 3.1 establishes that this is not possible in general if one only relies on ℓ_*.

Example 3.1. Consider a DSFM problem with a extremely sparse incidence structure with $|S_r| = 2$. More precisely, let $N = 2n+1$, $R = 2n$, and $\|\mu\|_1 = \sum_{r\in[R]} |S_r| = 4n \ll NR$. Let F_r be incident to the elements $\{r, r+1\}$, for all $r \in [R]$, and be such that $F_r(\{r\}) = F_r(\{r+1\}) = 1, F_r(\emptyset) = F_r(\{r, r+1\}) = 0$. Then, $\ell_* < \frac{7}{N^2}$.

Note that the optimal solution of problem (4) for this particular setting equals $y^* = 0$. Let us consider a point $y \in \mathcal{B}$ specified as follows. First, due to the given incidence relations, the block y_r has two components corresponding to the elements indexed by r and $r+1$. For any $r \in [R]$,

$$y_{r,r} = -y_{r,r+1} = \begin{cases} \frac{r}{n} & r \leq n, \\ \frac{2n+1-r}{n} & r \geq n+1. \end{cases} \tag{7}$$

Therefore, $g(y) = \frac{1}{n},\ \|y\|_2^2 > \frac{4}{3}n$, which results in $\ell_* < \frac{3}{2n^2} \leq \frac{7}{N^2}$ for all $n \geq 3$.

Example 3.1 only illustrates that an important parameter of CDMs cannot be improved using incidence information; but this does not necessarily imply that a sequential RCDM that uses incidence structures cannot offer better convergence rates than $O(N^2R)$. In Section E of the Supplement, we present additional experimental evidence that supports our observation, using the setting of Example 3.1.

As a final remark, note that Nishihara et al. [15] also proposed a lower bound that does not make use of sparse incidence structures and only works for the AP method.

3.3 New Parallel CD methods

In what follows, we propose two CDMs which rely on parallel projections and incidence relations.

The following observation is key to understanding the proposed approach. Suppose that we have a nonempty group of blocks $C \subseteq [R]$. Let $y, h \in \otimes_{r=1}^R \mathbb{R}^N$. If $h_{r,i}$ is nonzero only for block $r \in C$ and $i \in S_r$, then,

$$g(y+h) = g(y) + \langle \nabla g(y), h \rangle + \frac{1}{2}\|Ah\|_2^2 \leq\ g(y) + \sum_{r\in C} \langle \nabla_r g(y), h_r \rangle + \sum_{r \in C} \frac{1}{2}\|h_r\|_{2,\mu^C}^2. \tag{8}$$

Hence, for all $r \in C$, if we perform projections onto $\mathcal{B}_r$ with respect to the norm $\|\cdot\|_{2,\mu^C}$ simultaneously in each iteration of the CDM, convergence is guaranteed as the value of the objective function remains bounded. The smaller the components of μ^C, the faster the convergence. Note that the components of μ^C are the numbers of incidence relations of elements restricted to the set C. Hence, in each iteration, blocks that ought to be updated in parallel are those that correspond to submodular functions that have supports with smallest possible intersections.

One can select blocks that are to be updated in parallel in a combinatorially specified fashion or in a randomized fashion, as dictated by what we call an α-proper distribution. To describe our parallel RCDM, we first introduce the notion of an α-proper distribution.

Definition 3.8. Let P be a distribution used to sample a group of C blocks. Define $\theta^P = (\theta_1^P, \theta_2^P, ..., \theta_R^P)$ such that for $r \in [R]$, $\theta_r^P \triangleq \mathbb{E}_{C\sim P}\left[\mu^C | r \in C\right]$. We say that P is an α-proper distribution, if for any $r \in [R]$ and a given $\alpha \in (0,1)$, we have $\mathbb{P}(r \in C) = \alpha$.

We are now ready to describe the parallel RCDM algorithm – Algorithm 1; the description of the parallel ACDM is postponed to Section J of the Supplement.

Algorithm 1: Parallel RCDM for Solving (4)

Input: $\mathcal{B}$, α

0: Initialize $y^{(0)} \in \mathcal{B}$, $k \leftarrow 0$

1: Do the following steps iteratively until the dual gap $< \epsilon$:

2: Sample C_{i_k} using some α-proper distribution P

3: For $r \in C_{i_k}$:

4: $y_r^{(k+1)} \leftarrow \Pi_{\mathcal{B}_r, \theta_r^P}(y_r^{(k)} - (\theta_r^P)^{-1} \odot \nabla_r g(y^{(k)}))$

5: Set $y_r^{(k+1)} \leftarrow y_r^{(k)}$ for $r \notin C_{i_k}$, $k \leftarrow k+1$

6: Output $y^{(k)}$

Next, we establish strong convexity results for the space $\|\cdot\|_{2,\theta^P}$ by invoking Lemma 3.1.

Lemma 3.9. For any $y \in \mathcal{B}$, let $y^* = \arg\min_{\xi\in\Xi} \|\xi - y\|_{2,\theta^P}^2$. Then,

$$\|Ay - Ay^*\|_2^2 \geq \frac{2}{N\|\theta^P\|_{1,\infty}} \|y - y^*\|_{2,\theta^P}^2.$$

The convergence rate of Algorithm 1 is established in the next theorem.

Theorem 3.10. At each iteration of Algorithm 1, $y^{(k)}$ satisfies

$$\mathbb{E}\left[g(y^{(k)}) - g(y^*) + \frac{1}{2}d_{\theta^P}^2(y^k, \xi)\right] \leq \left[1 - \frac{4\alpha}{(N\|\theta^P\|_{1,\infty} + 2)}\right]^k \left[g(y^{(0)}) - g(y^*) + \frac{1}{2}d_{\theta^P}^2(y^0, \xi)\right].$$

The parameter $N\|\theta^P\|_{1,\infty}$ is obtained by combining the strong convexity constant and the properties of the sampling distribution P. Small values of $\|\theta^P\|_{1,\infty}$ ensure better convergence rates, and we next bound this value.

Lemma 3.11. For any α-proper distribution P and an element $i \in [N]$, $\max_{r\in[R]:i\in S_r} \theta_{r,i}^P \geq \max\{\alpha\mu_i, 1\}$. Consequently, $\|\theta^P\|_{1,\infty} \geq \max\{\alpha\|\mu\|_1, N\}$.

Without considering incidence relations, i.e., by setting $\|\mu\|_1 = NR$, one always has $\|\theta^P\|_{1,\infty} \geq \alpha NR$, which shows that parallelization cannot improve the convergence rate of the RCDM.

The next lemma characterizes an achievable $\|\theta^P\|_{1,\infty}$ obtained by choosing P to be a uniform distribution, which, when combined with Theorem 3.10, proves the result of the last column in Table 1.

Lemma 3.12. If C is a set of size $0 < K \leq R$ obtained by sampling the K-subsets of $[R]$ uniformly at random, then $\theta_r^P = \frac{K-1}{R-1}\mu + \frac{R-K}{R-1}1$. Moreover, $\|\theta^P\|_{1,\infty} = \frac{K-1}{R-1}\|\mu\|_1 + \frac{R-K}{R-1}N$.

Comparing Lemma 3.11 and Lemma 3.12, we see that the $\|\theta^P\|_{1,\infty}$ achieved by sampling uniformly at random is at most a factor of two of the lower bound since $\alpha = K/R$. A natural question is if it is possible to devise a better sampling strategy. This question is addressed in Section K of the Supplement, where we related the sampling problem to equitable coloring [20]. By using Hajnal-Szemerédi's Theorem [25], we derived a sufficient condition under which an α-proper distribution P that achieves the lower bound in Lemma 3.11 can be found in polynomial time. We also described a greedy algorithm for minimizing $\|\theta^P\|_{1,\infty}$ that empirically convergences faster than sampling uniformly at random.

4 Experiments

In what follows, we illustrate the performance of the newly proposed DSFM algorithms on a benchmark datasets used for MAP inference in image segmentation [9] and used for semi-supervised

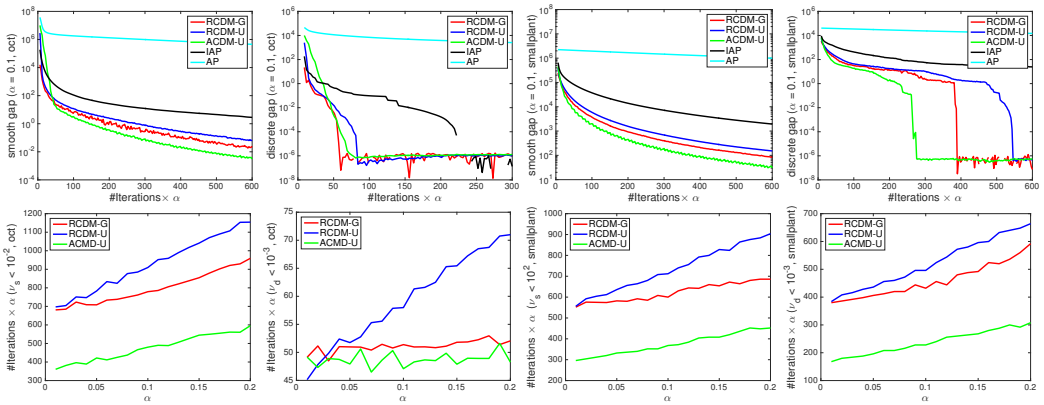

Figure 2: Image segmentation example. First row: Gap vs the number of iterations $\times\alpha$. Second row: The number of iterations $\times\alpha$ vs α. Here, α is the parallelization parameter, while $K = \alpha R$ equals the number of projections that have to be computed in each iteration.

learning over graphs [1]. More experiments on semi-supervised learning over hypergraphs can be found in Section M of the Supplement.

In all the experiments, we evaluated the convergence rate of the algorithms by using the smooth duality gap ν_s and the discrete duality gap ν_d. The primal problem solution equals $x = -Ay$ so that the smooth duality gap can be computed according to $\nu_s = \sum_r f_r(x) + \frac{1}{2}\|x\|^2 - (-\frac{1}{2}\|Ay\|^2)$. Moreover, as the level set $S_\lambda = \{v \in [N] | x_v > \lambda\}$ can be easily found based on x, the discrete duality gap can be written as $\nu_d = \min_\lambda F(S_\lambda) - \sum_{v\in[N]} \min\{-x_v, 0\}$.

MAP inference. We used two images – *oct* and *smallplant* – adopted from [14][2]. The images comprise 640×427 pixels so that $N = 273,280$. The decomposable submodular functions are constructed following a standard procedure. The first class of functions arises from the 4-neighbor grid graph over the pixels. Each edge corresponds to a pairwise potential between two adjacent pixels i, j that follows the formula $\exp(-\|v_i - v_j\|_2^2)$, where v_i is the RGB color vector of pixel i. We split the vertical and horizontal edges into rows and columns that result in $639 + 426 = 1065$ components in the decomposition. Note that within each row or each column, the edges have no overlapping pixels, so the projections of these submodular functions onto the base polytopes reduce to projections onto the base polytopes of edge-like submodular functions. The second class of submodular functions contain clique potentials corresponding to the superpixel regions; specifically, for region r, $F_r(S) = |S|(|S_r| - |S|)$ [26]. These functions give another 500 decomposition components. We apply the divide and conquer method in [14] to compute the projections required for this type of submodular functions. Note that in each experiment, all components of the submodular function are of nearly the same size, and thus the projections performed for different components incur similar computational costs. As the projections represent the primary computational units, for comparative purposes we use the number of iterations (similarly to [14, 16]).

We compared five algorithms: RCDM with a sampling distribution P found by the greedy algorithm (RCDM-G), RCDM with uniform sampling (RCDM-U), ACDM with uniform sampling (ACDM-U), AP based on (5) (IAP) and AP based on (3) (AP). Figure 2 depicts the results. In the first row, we compared the convergence rates of different algorithms for a fixed parallelization parameter $\alpha = 0.1$. The values on the horizontal axis correspond to # iterations $\times\,\alpha$, the total number of projections performed divided by R. The results are averaged over 10 independent experiments. We observe that the CD-based methods outperform AP-based methods, and that ACDM-U is the best performing CD-based method. IAP significantly outperforms AP. Similarly, RCDM-G outperforms RCDM-U. We also investigated the relationship between the number of iterations and the parameter α. We recorded the number of iterations needed to achieve a smooth and discrete gap below a certain given threshold. The results are shown in the second row of Figure 2. We did not plot the curves for the AP-based methods as they are essentially horizontal lines. Among the CD-based methods, ACDM-U performs best. RCDM-G offers a much better convergence rate than RCDM-U since the sampling probability P produced by the greedy algorithm leads to a smaller value of $\|\theta^P\|_{1,\infty}$ compared to

[1]The code for this work can be found in https://github.com/lipan00123/DSFM-with-incidence-relations.

[2]Downloaded from the website of Professor Stefanie Jegelka: http://people.csail.mit.edu/stefje/code.html

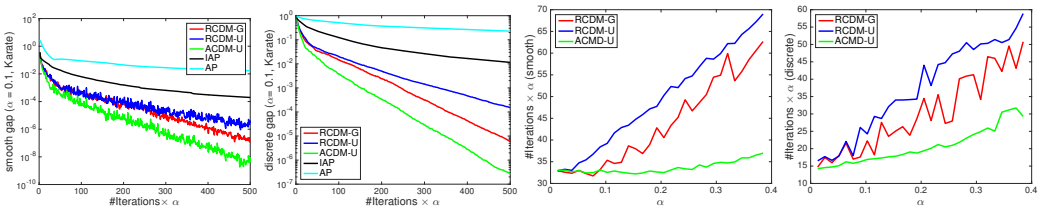

Figure 3: Zachary's Karate Club. Left two: Gap vs the number of iterations $\times\alpha$. Right two: The number of iterations $\times\alpha$ vs α. Here, α is the parallelization parameter, while $K = \alpha R$ equals the number of projections that have to be computed in each iteration.

uniform sampling. The reason behind this finding is that the supports of the components in the decomposition are localized, which makes the sampling P obtained from the greedy algorithm highly effective. For RCDM-U, the total number of iterations increases almost linearly with α ($= K/R$), which confirms the results of Lemma 3.12.

Note that in the above examples of MAP inference, another way to decompose the submodular functions is available: as there are three natural layers of non-overlapping incidence sets, we can merge all vertical edges, all horizontal edges, and all superpixel regions into three components respectively. Then, each of this component is incident to all pixels, and the derived results in this work will reduce to those of the former works [14, 16]. However, such a way to decompose submodular function strongly depends on the particular structure and thus is not general for DSFM problems. The following example on semi-supervised learning over graphs does not contain natural layers for decomposition.

Semi-supervised learning. We tested our algorithms over the dataset of Zachary's karate club [27]. This dataset is used as a benchmark example for evaluating semisupervised learning algorithms over graphs [28]. It includes $N = 34$ vertices and $R = 78$ submodular functions in the decomposition, each corresponding to one edge in the network. The objective function of both semi-supervised learning problems may be written as

$$\min_x \tau \sum_{r\in[R]} f_r(x) + \frac{1}{2}\|x - x_0\|_2^2 \tag{9}$$

where τ is a parameter that needs to be tuned, and $x_0 \in \{-1, 0, 1\}^N$, so that the nonzero components correspond to the labels that are known a priori. In our case, as we are only concerned with the convergence rate of the algorithm, we fix $\tau = 0.1$. In the experiments for Zachary's karate club, we set $x_0(1) = 1$, $x_0(34) = -1$ and let all other components of x_0 be equal to zero.

Figure 3 shows the results of the experiments pertaining to Zachary's karate club. In the left two subfigures, we compared the convergence rates of different algorithms for a fixed parallelization parameter $\alpha = 0.1$. The values on the horizontal axis correspond to # iterations $\times\ \alpha$, the total number of projections performed divided by R. In the right two subfigures, we controlled the numbers of projections executed within one iteration by tuning the parameter α and recorded the number of iterations needed to achieve smooth/discrete gaps below 10^{-3}. The values depicted on the vertical axis correspond to # iterations $\times\alpha$, describing the total number of projections needed to achieve the given accuracy. In all cases, we see the similar tendency to that of the MAP inference. As may be seen, AP-based methods require more projections than CD-based methods, but IAP consistently outperforms AP, which is consistent with our theoretical results. Among the CD-based methods, ACDM-U offers the best performance in general, and RCDM-G slightly outperforms RCDM-U, since the greedy algorithm used for sampling produces a smaller $\|\theta^P\|_{1,\infty}$ than uniform sampling. As the AP-based methods are completely parallelizable, and increasing the parameter α does not increase the total number of projections. However, for RCDM-U, the total number of iterations required increases almost linearly with α, which is supported by the result in Lemma 3.12. The performance curve for RCDM-G exhibits large oscillations due to the discrete problem component, needed for finding a balanced partition.

5 Acknowledgement

The authors gratefully acknowledge many useful suggestions by the reviewers. This work was supported in part by the NSF grant CCF 15-27636, the NSF Purdue 4101-38050 and the NFT STC center Science of Information.

References

[1] S. Fujishige, *Submodular functions and optimization*. Elsevier, 2005, vol. 58.

[2] K. Wei, R. Iyer, and J. Bilmes, "Submodularity in data subset selection and active learning," in *Proceedings of the International Conference on Machine Learning*, 2015, pp. 1954–1963.

[3] P. Li and O. Milenkovic, "Inhomogeneous hypergraph clustering with applications," in *Advances in Neural Information Processing Systems*, 2017, pp. 2305–2315.

[4] ——, "Submodular hypergraphs: p-laplacians, cheeger inequalities and spectral clustering," in *Proceedings of the International Conference on Machine Learning*, 2018, pp. 3014–3023.

[5] P. Kohli, P. H. Torr *et al.*, "Robust higher order potentials for enforcing label consistency," *International Journal of Computer Vision*, vol. 82, no. 3, pp. 302–324, 2009.

[6] H. Lin and J. Bilmes, "A class of submodular functions for document summarization," in *Proceedings of the Meeting of the Association for Computational Linguistics: Human Language Technologies-Volume 1*. Association for Computational Linguistics, 2011, pp. 510–520.

[7] A. Krause and C. Guestrin, "Near-optimal observation selection using submodular functions," in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 7, 2007, pp. 1650–1654.

[8] Y. T. Lee, A. Sidford, and S. C.-w. Wong, "A faster cutting plane method and its implications for combinatorial and convex optimization," in *Foundations of Computer Science (FOCS), 2015 IEEE 56th Annual Symposium on*. IEEE, 2015, pp. 1049–1065.

[9] P. Stobbe and A. Krause, "Efficient minimization of decomposable submodular functions," in *Advances in Neural Information Processing Systems*, 2010, pp. 2208–2216.

[10] V. Kolmogorov, "Minimizing a sum of submodular functions," *Discrete Applied Mathematics*, vol. 160, no. 15, pp. 2246–2258, 2012.

[11] A. Ene, H. Nguyen, and L. A. Végh, "Decomposable submodular function minimization: discrete and continuous," in *Advances in Neural Information Processing Systems*, 2017, pp. 2874–2884.

[12] F. Bach *et al.*, "Learning with submodular functions: A convex optimization perspective," *Foundations and Trends® in Machine Learning*, vol. 6, no. 2-3, pp. 145–373, 2013.

[13] L. Lovász, "Submodular functions and convexity," in *Mathematical Programming The State of the Art*. Springer, 1983, pp. 235–257.

[14] S. Jegelka, F. Bach, and S. Sra, "Reflection methods for user-friendly submodular optimization," in *Advances in Neural Information Processing Systems*, 2013, pp. 1313–1321.

[15] R. Nishihara, S. Jegelka, and M. I. Jordan, "On the convergence rate of decomposable submodular function minimization," in *Advances in Neural Information Processing Systems*, 2014, pp. 640–648.

[16] A. Ene and H. Nguyen, "Random coordinate descent methods for minimizing decomposable submodular functions," in *Proceedings of the International Conference on Machine Learning*, 2015, pp. 787–795.

[17] D. R. Karger, "Global min-cuts in RNC, and other ramifications of a simple min-cut algorithm." in *Proceedings of the ACM-SIAM Symposium on Discrete Algorithms*, vol. 93, 1993, pp. 21–30.

[18] C. Chekuri and C. Xu, "Computing minimum cuts in hypergraphs," in *Proceedings of the ACM-SIAM Symposium on Discrete Algorithms*. Society for Industrial and Applied Mathematics, 2017, pp. 1085–1100.

[19] J. Djolonga and A. Krause, "Scalable variational inference in log-supermodular models." in *Proceedings of the International Conference on Machine Learning*, 2015, pp. 1804–1813.

[20] W. Meyer, "Equitable coloring," *The American Mathematical Monthly*, vol. 80, no. 8, pp. 920–922, 1973.

[21] P. Wolfe, "Finding the nearest point in a polytope," *Mathematical Programming*, vol. 11, no. 1, pp. 128–149, 1976.

[22] D. Chakrabarty, P. Jain, and P. Kothari, "Provable submodular minimization using Wolfe's algorithm," in *Advances in Neural Information Processing Systems*, 2014, pp. 802–809.

[23] H. Karimi, J. Nutini, and M. Schmidt, "Linear convergence of gradient and proximal-gradient methods under the polyak-łojasiewicz condition," in *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*. Springer, 2016, pp. 795–811.

[24] Y. Nesterov, "Efficiency of coordinate descent methods on huge-scale optimization problems," *SIAM Journal on Optimization*, vol. 22, no. 2, pp. 341–362, 2012.

[25] A. Hajnal and E. Szemerédi, "Proof of a conjecture of Erdös," *Combinatorial Theory and Its Applications*, vol. 2, pp. 601–623, 1970.

[26] A. Levinshtein, A. Stere, K. N. Kutulakos, D. J. Fleet, S. J. Dickinson, and K. Siddiqi, "Turbopixels: fast superpixels using geometric flows," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 31, no. 12, pp. 2290–2297, 2009.

[27] W. W. Zachary, "An information flow model for conflict and fission in small groups," *Journal of Anthropological Research*, vol. 33, no. 4, pp. 452–473, 1977.

[28] T. N. Kipf and M. Welling, "Semi-supervised classification with graph convolutional networks," *arXiv preprint arXiv:1609.02907*, 2016.

[29] S. Fujishige and X. Zhang, "New algorithms for the intersection problem of submodular systems," *Japan Journal of Industrial and Applied Mathematics*, vol. 9, no. 3, p. 369, 1992.

[30] O. Fercoq and P. Richtárik, "Accelerated, parallel, and proximal coordinate descent," *SIAM Journal on Optimization*, vol. 25, no. 4, pp. 1997–2023, 2015.

[31] H. A. Kierstead, A. V. Kostochka, M. Mydlarz, and E. Szemerédi, "A fast algorithm for equitable coloring," *Combinatorica*, vol. 30, no. 2, pp. 217–224, 2010.

[32] A. Chambolle and J. Darbon, "On total variation minimization and surface evolution using parametric maximum flows," *International journal of computer vision*, vol. 84, no. 3, p. 288, 2009.

[33] R. Albert and A.-L. Barabási, "Statistical mechanics of complex networks," *Reviews of modern physics*, vol. 74, no. 1, p. 47, 2002.

[34] M. Hein, S. Setzer, L. Jost, and S. S. Rangapuram, "The total variation on hypergraphs-learning on hypergraphs revisited," in *Advances in Neural Information Processing Systems*, 2013, pp. 2427–2435.

[35] N. Yadati, M. Nimishakavi, P. Yadav, A. Louis, and P. Talukdar, "Hypergcn: Hypergraph convolutional networks for semi-supervised classification," *arXiv preprint arXiv:1809.02589*, 2018.

Supplement

A Discrete Optimization Approach for Computing the Projections $\Pi_{\mathcal{B}_r,w}(\cdot)$

The following Lemma A.1 describes how the projections $\Pi_{\mathcal{B}_r,w}(\cdot)$ can be performed via discrete optimization. Discrete methods are especially useful when $F_r(S)$ is concave in $|S|$, as in this case they have much smaller complexity than the min-norm algorithm of Wolfe [21].

Lemma A.1. The optimization problem $\min_{y_r\in\mathcal{B}_r}\|z-y_r\|_{2,w}^2$ is the dual of the problem $\min_{x\in\mathbb{R}^N} f_r(x)-\langle x,z\rangle+\frac{1}{2}\|x\|_{2,w^{-1}}^2$. A solution with coordinate accuracy ϵ for the latter setting can be obtained by solving the discrete problem

$$\min_S F_r(S)-z(S)+\lambda\sum_{i\in S_r\cap S} w_i^{-1},$$

where

$$\lambda\in\left[\min_{i\in[N]}[-F_r(\{i\})+z(\{i\})]w_i,\ \max_{i\in[N]}[F_r([N]/\{i\})-F_r([N])+z(\{i\})]w_i\right],$$

at most $\min\{|S_r|,\log 1/\epsilon\}$ times. The parameter λ is chosen based on a binary search procedure which requires solving the discrete problem $O(\log 1/\epsilon)$ times.

Proof. The first statement follows from $f_r(x)=\max_{y_r\in\mathcal{B}_r}\langle y_r,x\rangle$ and some simple algebra. The second claim follows from the divide and conquer algorithm described in Appendix B of [14]. □

B Proof of Lemma 3.1

The first part of the proof follows along the same line as the corresponding proof of Ene et al. [11] which is based on a submodular auxiliary graph and the path-augmentation algorithm [29], described in what follows.

Let $G=(V,E)$ be a directed graph such that the vertex set V corresponds to the elements in $[N]$, and where the arc set may be written as $E=\cup_{r\in[R]}E_r$, with E_r corresponding to a complete directed graph on the set of elements S_r incident to F_r. With each arc (u,v), we associate a capacity value based on a $y'\in\mathcal{B}$ according to $c(u,v)\triangleq\min\{f_r(S)-y'_r(S):S\subseteq S_r,u\in S,v\notin S\}$.

Next, we consider a procedure termed path augmentations over G that sequentially transforms y' from $y'=y$ to a point in $\mathcal{B}$ that satisfies $Ay'=z$; the vector y' is kept within $\mathcal{B}$ during the whole procedure. Let the set of source and sink nodes of the graph be defined as $N\triangleq\{v\in[N]|(Ay')_v<z_v\}$ and $P\triangleq\{v\in[N]|(Ay')_v>z_v\}$, where z is as defined in the statement of the lemma. If $N=P=\emptyset$, we have $Ay'=z$. It can be shown that there always exists a directed path with positive capacity from N to P unless $N=P=\emptyset$ [11]. In each step, we find the shortest directed path, denoted by $\mathcal{Q}$, with positive capacity from N to P. For each arc (u,v) in $\mathcal{Q}$, if the arc belongs to E_r, we set $y'_{r,u}\leftarrow y'_{r,u}+\rho$, $y'_{r,v}\leftarrow y'_{r,v}-\rho$, where ρ denotes the smallest capacity of any arc in $\mathcal{Q}$. This procedure ensures that $y'\in\mathcal{B}$ and that the procedure terminates in a finite number of steps, with $N=P=\emptyset$ [29].

The second part of the proof differs from the derivations of Ene et al. [11]. Suppose that $\{y'^{(0)}=y,y'^{(1)},...,y'^{(t)}\}$ is a sequence such that $y'^{(i)}$ equals the vector y' after the i-th step of the above procedure. We also assume that $Ay'^{(t)}=z$, implying that the algorithm terminated at step t. Hence, the point $y'^{(t)}$ is the desired value of ξ. During path-augmentation, no element appears in more than two updated arcs. Hence,

$$\|y'^{(i)}-y'^{(i-1)}\|_{2,\theta}\leq\sqrt{2\sum_v\max_{r\in[R]:v\in S_r}\theta_{r,v}}\rho=\sqrt{2\|\theta\|_{1,\infty}}\rho.$$

As $\|Ay'^{(i)}-Ay'^{(i-1)}\|_1=2\rho$, we have

$$\|y'^{(i)}-y'^{(i-1)}\|_{2,\theta}\leq\sqrt{\frac{\|\theta\|_{1,\infty}}{2}}\|Ay'^{(i)}-Ay'^{(i-1)}\|_1.$$

An important observation is that during the path-augmentation procedure, for each component $v \in [N]$, the updated sequence $\{(Ay'^{(i)})_v\}_{i=1,2,..,t}$ converges monotonically to z_v. Hence, $\|Ay'^{(t)} - Ay'^{(0)}\|_1 = \sum_{i=1}^{t} \|Ay'^{(i)} - Ay'^{(i-1)}\|_1$. By using the triangle inequality for the norm $\|\cdot\|_{2,\theta}$, we obtain

$$\sqrt{\frac{\|\theta\|_{1,\infty}}{2}}\|z - Ay\|_1 = \sqrt{\frac{\|\theta\|_{1,\infty}}{2}}\|Ay'^{(t)} - Ay'^{(0)}\|_1 \geq \sum_{i=1}^{t} \|y'^{(i)} - y'^{(i-1)}\|_{2,\theta}$$
$$\geq \|y'^{(t)} - y'^{(0)}\|_{2,\theta} = \|y'^{(t)} - y\|_{2,\theta}.$$

Invoking the Cauchy-Schwarz inequality establishes $\|z - Ay\|_1 \leq \sqrt{\|w^{-1}\|_1}\|z - Ay\|_{2,w}$, which concludes the proof.

C Proof for Lemma 3.2

The equivalence between problem (5) and problem (4) is easy to establish, as y is obtained from y' by simply removing its zero components. The second statement is proved as follows:

$$\begin{aligned}
&\min_{y\in\mathcal{B}} \min_{a:Aa=0, a_{r,i}=0,\ \forall(r,i):i\notin S_r} \frac{1}{2}\|y-a\|^2_{2,I(\mu)} \\
=&\min_{y\in\mathcal{B}} \min_{a:a_{r,i}=0,\ \forall(r,i):i\notin S_r} \max_{\lambda\in\mathbb{R}^N} \frac{1}{2}\|y-a\|^2_{2,I(\mu)} - \langle\lambda, Aa\rangle \\
\overset{1)}{=}&\min_{y\in\mathcal{B}} \max_{\lambda\in\mathbb{R}^N} \min_{a\in\otimes_{r=1}^{R}R^N} \frac{1}{2}\sum_{r\in[R]}\sum_{i\in S_r}[\mu_i(y_{r,i}-a_{r,i})^2 - 2\lambda_i a_{r,i}] \\
=&\min_{y\in\mathcal{B}} \max_{\lambda\in\mathbb{R}^N} \frac{1}{2}\sum_{r\in[R]}\sum_{i\in S_r}[\mu_i^{-1}\lambda_i^2 - 2\lambda_i(\mu_i^{-1}\lambda_i + y_{r,i})] \\
=&\min_{y\in\mathcal{B}} \max_{\lambda\in\mathbb{R}^N} \frac{1}{2}\sum_{r\in[R]}\sum_{i\in S_r}(-\mu_i^{-1}\lambda_i^2 - 2\lambda_i y_{r,i}) \\
\overset{2)}{=}&\min_{y\in\mathcal{B}} \max_{\lambda\in\mathbb{R}^N} -\frac{1}{2}\|\lambda\|_2^2 - \langle\lambda, Ay\rangle \\
=&\min_{y\in\mathcal{B}} \|Ay\|_2^2,
\end{aligned}$$

where 1) is obtained using the incidence relations $y_{r,i} = a_{r,i} = 0$ if $i \notin S_r$ and 2) holds because $\mu_i = |\{r \in [R] | i \in S_r\}|$. The optimal y, a, λ satisfy $a_{r,i} = y_{r,i} + \mu_i^{-1}\lambda_i$ for all $i \in S_r$, $r \in [R]$ and $\lambda = -Ay$.

D Proof of Lemma 3.4

First, consider a $y \in \mathcal{B}/\Xi$. We have $d_{I(\mu)}(y, \mathcal{Z}) = \|Ay + x^*\|_2$, since

$$\begin{aligned}
\frac{1}{2}d_{I(\mu)}(y,\mathcal{Z})^2 &= \min_{a\in\mathcal{Z}} \frac{1}{2}\|y-a\|^2_{2,I(\mu)} \\
&= \min_{a:a_{r,i}=0,\ \forall(r,i):i\notin S_r} \max_{\lambda\in\mathbb{R}^N} \frac{1}{2}\|y-a\|^2_{2,I(\mu)} - \langle\lambda, Aa + x^*\rangle \\
&\overset{1)}{=} \max_{\lambda\in\mathbb{R}^N} \min_{a\in\otimes_{r=1}^{R}\mathbb{R}^N} \frac{1}{2}\sum_{r\in[R]}\sum_{i\in S_r}[\mu_i(y_{r,i}-a_{r,i})^2 - 2\lambda_i a_{r,i}] - \langle\lambda, x^*\rangle \\
&= \max_{\lambda\in\mathbb{R}^N} \frac{1}{2}\sum_{r\in[R]}\sum_{i\in S_r}[-\mu_i^{-1}\lambda_i^2 - \lambda_i y_{r,i}] - \langle\lambda, x^*\rangle \\
&\overset{2)}{=} \max_{\lambda\in\mathbb{R}^N} -\frac{1}{2}\|\lambda\|_2^2 - \lambda^T(Ay + x^*) \\
&= \frac{1}{2}\|Ay + x^*\|_2^2.
\end{aligned}$$

where 1) is obtained using the incidence relations $y_{r,i} = a_{r,i} = 0$ if $i \notin S_r$ and 2) holds because $\mu_i = |\{r \in [R] | i \in S_r\}|$. Based on Lemma 3.1, we know that there exists a $\xi \in \mathcal{B}$ such that $A\xi = -x^*$ and

$$\|y - \xi\|_{2,I(\mu)} \leq \sqrt{\frac{N\|I(\mu)\|_{1,\infty}}{2}} \|Ay + x^*\|_2 = \sqrt{\frac{N\|\mu\|_1}{2}} \|Ay + x^*\|_2.$$

Therefore, $\kappa(y) = \frac{d_{I(\mu)}(y,\Xi)}{d_{I(\mu)}(y,\mathcal{Z})} \leq \sqrt{\frac{N\|\mu\|_1}{2}}$.

Next, consider a $y \in \mathcal{Z}/\Xi$. As $\mathcal{B}$ is compact, there exists a $y' \in \mathcal{B}$ that achieves $d_{I(\mu)}(y, \mathcal{B}) = \|y - y'\|_{2,I(\mu)}$. Based on Lemma 3.1, we also know that there exists a $\xi \in \mathcal{B}$ such that $A\xi = -x^*$ and

$$\|\xi - y'\|_{2,I(\mu)} \leq \sqrt{\frac{\|I(\mu)\|_{1,\infty}}{2}} \|Ay' + x^*\|_1 = \sqrt{\frac{\|\mu\|_1}{2}} \|Ay' + x^*\|_1.$$

Moreover, we have

$$\begin{aligned}\|Ay' + x^*\|_1 =& \|Ay' - Ay\|_1 \leq \|y' - y\|_1 = \sum_{v\in[N]} \sum_{r:v\in S_r} |y'_{r,v} - y_{r,v}| \\ &\leq \sum_{v\in[N]} \left[\mu_v \sum_{r:v\in S_r} (y'_{r,v} - y_{r,v})^2\right]^{\frac{1}{2}} \leq \sqrt{N}\|y' - y\|_{2,I(\mu)}.\end{aligned}$$

As $\xi \in \Xi$, it holds that $d_{I(\mu)}(y, \Xi) \leq \|\xi - y\|_{2,I(\mu)} \leq \|y' - y\|_{2,I(\mu)} + \|y' - \xi\|_{2,I(\mu)}$. In addition, as

$$\|y' - \xi\|_{2,I(\mu)} \leq \sqrt{\frac{\delta^s}{2}} \|Ay' + x^*\|_1 \leq \sqrt{\frac{N\|\mu\|_1}{2}} \|y' - y\|_{2,I(\mu)},$$

we know that $d_{I(\mu)}(y, \Xi) \leq (1 + \sqrt{\frac{N\|\mu\|_1}{2}})\|y' - y\|_{2,I(\mu)}$. Therefore,

$$\kappa(y) = \frac{d_{I(\mu)}(y, \Xi)}{d_{I(\mu)}(y, \mathcal{B})} \leq \left(1 + \sqrt{\frac{N\|\mu\|_1}{2}}\right),$$

which concludes the proof.

E Simulation for Example 3.1

We provide additional empirical evidence that the convergence result suggested by the bound on $\ell_* \leq \frac{7}{N^2}$ is correct. We constructed a DSFM problem following Example 3.1 and initialized y according to equation (7). We used the number of iterations k required to attain $g(y^{(k)}) \leq \epsilon g(y^{(0)})$ as a measure for the speed of convergence. We ran the simulations for $n \in [5, 50]$ and averaged the results for each n over 10 independent runs. Figure 4 shows the results. The values next to the curves are their slopes obtained via a linear regression involving $\ln(\#\text{ Iterations}) \sim \ln(N)$. As the accuracy threshold increases, the slope approaches the value 3, which indicates that the required number of iterations equals $O(N^2 R)$.

F Proof of Lemma 3.9

Choose $z = Ay^*$ in Lemma 3.1. Then, there is a $\xi \in \mathcal{B}$ such that $\|Ay - Ay^*\|^2 \geq \frac{2}{N\|\theta^P\|_{1,\infty}} \|y - \xi\|^2_{2,\theta^P}$. Moreover as $A\xi = z = Ay^* = -x^*$, we also have $\xi \in \Xi$. Therefore, $\|y - \xi\|^2_{2,\theta^P} \geq \|y - y^*\|^2_{2,\theta^P}$, which concludes the proof.

G Proof for Theorem 3.10

First, given a group of blocks C and $y \in \otimes_{r=1}^R \mathbb{R}^N$, we define $y_{[C]} \in \otimes_{r=1}^R \mathbb{R}^N$ as

$$(y_{[C]})_r = \begin{cases} y_r & \text{if } r \in C, \\ 0 & \text{if } r \notin C. \end{cases}$$

The following lemma holds.

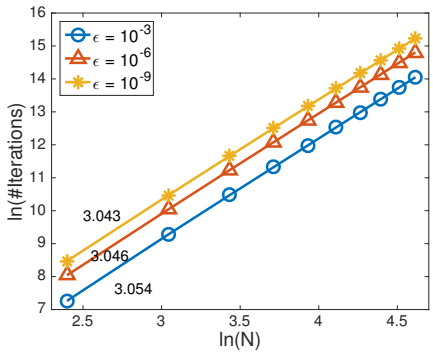

Figure 4: Simulations accompanying Example 3.1: $\ln$(the number of iterations) vs $\ln(N)$.

Lemma G.1. Let C be a group of blocks sampled according to a $\alpha-$proper distribution P. Then, for any $y \in \otimes_{r=1}^{R}\mathbb{R}^N$ and $y_{r,i} = 0$, whenever $i \notin S_r$, one has

$$\mathbb{E}_{C\sim P}(\|y_{[C]}\|^2_{2,I(\mu^C)}) = \mathbb{E}_{C\sim P}(\|y_{[C]}\|^2_{2,\theta^P}).$$

Proof.

$$\mathbb{E}_{C\sim P}(\|y_{[C]}\|^2_{2,I(\mu^C)}) = \mathbb{E}_{C\sim P}(\sum_{r\in C}\|y_r\|^2_{2,\mu^C}) = \sum_{r\in[R]}\mathbb{E}_{C\sim P}\left[\|y_r\|^2_{2,\mu^C}1_{r\in C}\right]$$

$$= \sum_{r\in[R]}\mathbb{E}\left[1_{r\in C}\mathbb{E}_{C\sim P}\left[\|y_r\|^2_{2,\mu^C}|r\in C\right]\right] = \sum_{r\in[R]}\mathbb{E}\left[1_{r\in C}\|y_r\|^2_{2,\theta^P_r}\right] = \mathbb{E}_{C\sim P}(\|y_{[C]}\|^2_{2,\theta^P}).$$

□

Next, we turn our attention to the proof of the theorem. For this purpose, suppose that $y^* = \arg\min_{y\in\Xi}\|y - y^{(k)}\|_{2,\theta^P}$.

G.1 Algorithm 1

We start with by establishing the following results.

Lemma G.2. It can be shown that

$$\langle\nabla g(y^{(k)}), y^* - y^{(k)}\rangle \overset{1)}{\leq} g(y^*) - g(y^{(k)}) - \frac{1}{N\|\theta^P\|_{1,\infty}}\|y^{(k)} - y^*\|^2_{2,\theta^P}$$

$$\overset{2)}{\leq} \frac{4}{N\|\theta^P\|_{1,\infty}+2}\left[g(y^*) - g(y^{(k)}) - \frac{1}{2}\|y^{(k)} - y^*\|^2_{2,\theta^P}\right]. \tag{10}$$

Proof. From Lemma 3.8 we can infer that

$$\|Ay^{(k)} - Ay^*\|_2^2 \geq \frac{2}{N\|\theta^P\|_{1,\infty}}\|y^{(k)} - y^*\|^2_{2,\theta^P} \Rightarrow$$

$$g(y^*) \geq g(y^{(k)}) + \langle\nabla g(y^{(k)}), y^* - y^{(k)}\rangle + \frac{1}{N\|\theta^P\|_{1,\infty}}\|y^{(k)} - y^*\|^2_{2,\theta^P}, \tag{11}$$

$$g(y^{(k)}) \geq g(y^*) + \langle\nabla g(y^*), y^{(k)} - y^*\rangle + \frac{1}{N\|\theta^P\|_{1,\infty}}\|y^{(k)} - y^*\|^2_{2,\theta^P}. \tag{12}$$

As $\langle\nabla g(y^*), y^{(k)} - y^*\rangle \geq 0$, (12) gives

$$g(y^*) - g(y^{(k)}) \leq -\frac{1}{N\|\theta^P\|_{1,\infty}}\|y^{(k)} - y^*\|^2_{2,\theta^P}. \tag{13}$$

The inequality (11) yields claim 1) in (10). Claim 2) in (10) follows from (13). □

The following lemma is a direct consequence of the optimality of $y_r^{(k+1)}$ for an oblique projection.

Lemma G.3.

$$\langle \nabla_r g(y^{(k)}), y_r^{(k+1)} - y_r^* \rangle \leq \langle y_r^{(k)} - y_r^{(k+1)}, y_r^{(k+1)} - y_r^* \rangle_{\theta_r^P}.$$

The following lemma follows from a simple manipulation of the Euclidean norm.

Lemma G.4.

$$\begin{aligned}\frac{1}{2}\|y_r^{(k+1)} - y_r^{(k)}\|_{2,\theta_r^P}^2 &= \frac{1}{2}\|y_r^{(k+1)} - y_r^*\|_{2,\theta_r^P}^2 + \frac{1}{2}\|y_r^* - y_r^{(k)}\|_{2,\theta_r^P}^2 + \langle y_r^{(k+1)} - y_r^*, y_r^* - y_r^{(k)} \rangle_{\theta_r^P} \\ &= -\frac{1}{2}\|y_r^{(k+1)} - y_r^*\|_{2,\theta_r^P}^2 + \frac{1}{2}\|y_r^* - y_r^{(k)}\|_{2,\theta_r^P}^2 + \langle y_r^{(k+1)} - y_r^*, y_r^{(k+1)} - y_r^{(k)} \rangle_{\theta_r^P}\end{aligned}$$

Let us analyze next the amount by which the objective function decreases in each iteration. The following expectation is with respect to $C_{i_k} \sim P$.

$$\mathbb{E}\left[g(y^{(k+1)})\right] \tag{14}$$

$$\begin{aligned}&\overset{1)}{\leq} g(y^{(k)}) + \mathbb{E}\left\{\sum_{r \in C_{i_k}} \left[\langle \nabla_r g(y^{(k)}), y_r^{(k+1)} - y_r^{(k)} \rangle + \frac{1}{2}\|y_r^{(k+1)} - y_r^{(k)}\|_{2,\mu_r^{C_{i_k}}}^2\right]\right\} \\ &\overset{2)}{=} g(y^{(k)}) + \mathbb{E}\left\{\sum_{r \in C_{i_k}} \left[\langle \nabla_r g(y^{(k)}), y_r^{(k+1)} - y_r^{(k)} \rangle + \frac{1}{2}\|y_r^{(k+1)} - y_r^{(k)}\|_{2,\theta_r^P}^2\right]\right\} \\ &= g(y^{(k)}) + \mathbb{E}\left\{\sum_{r \in C_{i_k}} \left[\langle \nabla_r g(y^{(k)}), y_r^* - y_r^{(k)} \rangle + \langle \nabla_r g(y^{(k)}), y_r^{(k+1)} - y_r^* \rangle + \frac{1}{2}\|y_r^{(k+1)} - y_r^{(k)}\|_{2,\theta_r^P}^2\right]\right\} \\ &\overset{3)}{\leq} g(y^{(k)}) + \mathbb{E}\left\{\sum_{r \in C_{i_k}} \left[\langle \nabla_r g(y^{(k)}), y_r^* - y_r^{(k)} \rangle - \frac{1}{2}\|y_r^{(k+1)} - y_r^*\|_{2,\theta_r^P}^2 + \frac{1}{2}\|y_r^* - y_r^{(k)}\|_{2,\theta_r^P}^2\right]\right\} \\ &= g(y^{(k)}) + \alpha\langle \nabla g(y^{(k)}), y^* - y^{(k)} \rangle - \mathbb{E}\left[\frac{1}{2}\|y_{[C_{i_k}]}^{(k+1)} - y_{[C_{i_k}]}^*\|_{2,\theta^P}^2\right] + \mathbb{E}\left[\frac{1}{2}\|y_{[C_{i_k}]}^{(k)} - y_{[C_{i_k}]}^*\|_{2,\theta^P}^2\right] \\ &\overset{4)}{=} g(y^{(k)}) + \alpha\langle \nabla g(y^{(k)}), y^* - y^{(k)} \rangle - \mathbb{E}\left[\frac{1}{2}\|y^{(k+1)} - y^*\|_{2,\theta^P}^2\right] + \mathbb{E}\left[\frac{1}{2}\|y^{(k)} - y^*\|_{2,\theta^P}^2\right] \\ &\overset{5)}{\leq} g(y^*) - \mathbb{E}\left[\frac{1}{2}\|y^{(k+1)} - y^*\|_{2,\theta^P}^2\right] + \left[1 - \frac{4\alpha}{N\|\theta^P\|_{1,\infty} + 2}\right]\left\{g(y^{(k)}) - g(y^*) - \frac{1}{2}\|y^{(k)} - y^*\|_{2,\theta^P}^2\right\},\end{aligned} \tag{15}$$

where $1)$ follows from inequality (8), $2)$ holds due to Lemma G.1, $3)$ is a consequence of Lemma G.3 and Lemma G.4, $4)$ is due to $y_r^{(k+1)} = y_r^{(k)}$ for $r \notin C_{i_k}$, and $5)$ may be established from (10).

Equation (15) further establishes that

$$\begin{aligned}\mathbb{E}\left[g(y^{(k+1)}) - g(y^*) + \frac{1}{2}d_{\theta^P}^2(y^{k+1}, \xi)\right] &\leq \mathbb{E}\left[g(y^{(k+1)}) - g(y^*) + \frac{1}{2}\|y^{(k+1)} - y^*\|_{2,\theta^P}^2\right] \\ &\leq \left[1 - \frac{4\alpha}{N\|\theta^P\|_{1,\infty} + 2}\right]\mathbb{E}\left[g(y^{(k)}) - g(y^*) + \frac{1}{2}d_{\theta^P}^2(y^k, \xi)\right].\end{aligned}$$

The proof follows by repeating the derivations for all k.

H Proof of Lemma 3.11

According to the definition of θ^P, we have

$$\begin{aligned}
\max_{r\in[R]:i\in S_r}\theta^P_{r,i} &= \max_{r\in[R]:i\in S_r}\mathbb{E}_{C\sim P}\left[\mu_i^C|r\in C\right]\\
&= \max_{r\in[R]:i\in S_r}\mathbb{E}_{C\sim P}\left[\sum_{r'\in[R]:i\in S_{r'}}1_{r'\in C}|r\in C\right]\\
&= \max_{r\in[R]:i\in S_r}\sum_{r'\in[R]:i\in S_{r'}}\mathbb{P}_{C\sim P}\left[r'\in C|r\in C\right] \qquad (16)\\
&= \frac{1}{\alpha}\max_{r\in[R]:i\in S_r}\sum_{r'\in[R]:i\in S_{r'}}\mathbb{P}_{C\sim P}\left[r'\in C, r\in C\right]\\
&\geq \frac{1}{\alpha\mu_i}\sum_{r,r'\in[R]:i\in S_r,S_{r'}}\mathbb{P}_{C\sim P}\left[r'\in C, r\in C\right]\\
&= \frac{1}{\alpha\mu_i}\mathbb{E}_{C\sim P}\left[|\{(r,r')\in C\times C : i\in S_r, i\in S_{r'}\}|\right]\\
&= \frac{1}{\alpha\mu_i}\mathbb{E}_{C\sim P}\left[(\mu_i^C)^2\right] \geq \frac{1}{\alpha d_i}\left[\mathbb{E}_{C\sim P}(\mu_i^C)\right]^2 = \frac{1}{\alpha d_i}\left(\sum_C\sum_{r\in[R]:i\in S_r}1_{r\in C}\mathbb{P}(C)\right)^2\\
&= \frac{1}{\alpha\mu_i}\left(\sum_{r\in[R]:i\in S_r}\mathbb{P}_{C\sim P}[r\in C]\right)^2\\
&= \frac{1}{\alpha\mu_i}\left(\alpha\mu_i\right)^2 = \alpha\mu_i.
\end{aligned}$$

From (16), we also have $\sum_{r'\in[R]:i\in S_{r'}}\mathbb{P}_{C\sim P}\left[r'\in C|r\in C\right]\geq\mathbb{P}_{C\sim P}\left[r\in C|r\in C\right]=1$, which proves the claimed result.

I Proof of Lemma 3.12

Similar to what was established for (16), one can show that $\theta^P_{r,i} = \sum_{r'\in[R]:i\in S_{r'}}\mathbb{P}_{C\sim P}\left[r'\in C|r\in C\right]$.

Consider next the right hand side of this equation for $\alpha=\frac{K}{R}$. In this case, for some r and some $i\in S_r$, we have

$$\begin{aligned}
\sum_{r'\in[R]:i\in S_{r'}}\mathbb{P}_{C\sim P}\left[r'\in C|r\in C\right] &= \mathbb{P}_{C\sim P}\left[r\in C|r\in C\right] + \sum_{r'\in[R]:i\in S_{r'},r'\neq r}\mathbb{P}_{C\sim P}\left[r'\in C|r\in C\right]\\
&= 1+\frac{R}{K}\sum_{r':i\in S_{r'},r'\neq r}\mathbb{P}_{C\sim P}\left[r'\in C, r\in C\right]\\
&= 1+\frac{R}{K}(\mu_i-1)\frac{\binom{R-2}{K-2}}{\binom{R}{K}} = 1+\frac{K-1}{R-1}(\mu_i-1).
\end{aligned}$$

Therefore, $\theta^P_{r,i}=\frac{K-1}{R-1}\mu_i+\frac{R-K}{R-1}$ when P is a uniform distribution.

J Analysis of the Accelerated Coordinate Descend Method

In the ACDM setting, we used the APPROX framework proposed by Fercoq and Richtárik in [30] and adapted it to this particular problem. In the general APPROX framework, the norm in (8) is chosen as follows: consider an arbitrary function ϕ with the component-wise smoothness and strong

convexity property. For block r, one has $|\nabla_r \phi(x) - \nabla_r \phi(y)| \leq L_r \|x_r - y_r\|_{\nu_r}$, where $\|\cdot\|_{\nu_r}$ is a norm associated with the r-th block. If one wants to process multiple blocks simultaneously, say those in a group C, one first needs to find a constant L_C such that for any h as defined in (8), it holds that

$$\phi(y+h) \;\leq\; \phi(y) + \sum_{r\in C} \langle \nabla_r \phi(y), h_r \rangle + \sum_{r\in C} \frac{L_C}{2} \|h_r\|^2_{2,\nu_r}.$$

The smaller the value of the multiplier L_C, the faster the convergence. Typically, L_C lies in $[\max_{r\in C} L_r, \sum_{r\in C} L_r]$.

Recall the smoothness property of g from equation (6). A direct application of APPROX to our problem gives

$$g(y+h) \;\leq\; g(y) + \sum_{r\in C} \langle \nabla_r g(y), h_r \rangle + \sum_{r\in C} \frac{\max_{i\in[N]} \mu_i^C}{2} \|h_r\|_2^2.$$

As $(\max_{i\in[N]} \mu_i^C) \geq \mu_j^C$ for all $j \in [N]$, we obtain convergence rates worse than those implied by inequality (8). To actually obtain the guarantees in (8), one needs to dispose with the $\|\cdot\|_2$ norm at the block level and break the blocks into components corresponding to the individual elements. The elements are evaluated independently through the use of the norm $\|\cdot\|_{2,\mu^C}$.

Algorithm 2: Parallel ACDM for Solving (4)

Input: $\mathcal{B}$, α, some constant $c > 0$

0: Initialize $y^{(0)} \in \mathcal{B}$, $k \leftarrow 0$

1: $c' \leftarrow \left\lceil (1+c)\frac{\sqrt{2N\|\theta^P\|_{1,\infty}}}{\alpha} + c \right\rceil$

2: Do the following steps iteratively until the dual gap $< \epsilon$:

3: If $k = lc'$ for some $l \in \mathbb{Z}$, $z^{(k)} \leftarrow y^{(k)}$, $\lambda_k \leftarrow 1$

4: $p^{(k)} \leftarrow (1-\lambda_k) y^{(k)} + \lambda_k z^{(k)}$

5: Sample C_{i_k} using some α-proper distribution P

6: $z^{(k+1)} \leftarrow z^{(k)}$

7: For $r \in C_{i_k}$:

8: $z_r^{(k+1)} \leftarrow \Pi_{\mathcal{B}_r, \theta_r^P}(z_r^{(k)} - \frac{\alpha}{\lambda_k}(\theta_r^P)^{-1} \odot \nabla_r g(p^{(k)}))$

9: $y^{(k+1)} \leftarrow p^{(k)} + \frac{\lambda_k}{\alpha}(z^{(k+1)} - z^{(k)})$

10: $\lambda_{k+1} \leftarrow \frac{\sqrt{\lambda_k^4 + 4\lambda_k^2} - \lambda_k^2}{2}$

11: $k \leftarrow k+1$

12: Output $y^{(k)}$

Similar to the APPROX method [30], the parallel ACDM can also be implemented to avoid full-dimensional vector operations (see Section J.2). The following theorem characterizes the convergence property of Algorithm 2.

Theorem J.1. Given $c > 0$, let $c' = \left\lceil (1+c)\frac{\sqrt{2N\|\theta^P\|_{1,\infty}}}{\alpha} + c \right\rceil$. Consider the iterations $k = lc'$ for $l \in \mathbb{Z}_{\geq 0}$. Then, $y^{(k)}$ of Algorithm 2 satisfies

$$\mathbb{E}\left[g(y^{(k)} - g(y^*)\right] \leq \frac{1}{(1+c)^l}\left[g(y^{(0)}) - g(y^*)\right].$$

J.1 Proof of Theorem J.1

We start by establishing a number of background results.

The following lemma is due to the optimality of $z_r^{(k+1)}$.

Lemma J.2.

$$\langle \nabla_r g(p^{(k)}), z_r^{(k+1)} - y_r^* \rangle \leq \frac{\lambda_k}{\alpha} \langle z_r^{(k)} - z_r^{(k+1)}, z_r^{(k+1)} - y_r^* \rangle_{\theta_r^P}.$$

Once again, one can easily establish the following result pertaining to the Euclidean norm.

Lemma J.3.

$$\begin{aligned}\frac{1}{2}\|z_r^{(k+1)}-z_r^{(k)}\|_{2,\theta_r^P}^2 &= \frac{1}{2}\|z_r^{(k+1)}-y_r^*\|_{2,\theta_r^P}^2+\frac{1}{2}\|y_r^*-z_r^{(k)}\|_{2,\theta_r^P}^2+\langle z_r^{(k+1)}-y_r^*, y_r^*-z_r^{(k)}\rangle_{\theta_r^P}\\ &= -\frac{1}{2}\|z_r^{(k+1)}-y_r^*\|_{2,\theta_r^P}^2+\frac{1}{2}\|y_r^*-z_r^{(k)}\|_{2,\theta_r^P}^2+\langle z_r^{(k+1)}-y_r^*, z_r^{(k+1)}-z_r^{(k)}\rangle_{\theta_r^P}.\end{aligned}$$

The next result follows from the convexity property of the function g.

Lemma J.4.

$$\begin{aligned}\lambda_k\langle\nabla g(p^{(k)}), y^*-z^{(k)}\rangle &= \langle\nabla g(p^{(k)}), \lambda_k y^*-\lambda_k z^{(k)}\rangle = \langle\nabla g(p^{(k)}), \lambda_k y^*-(p^{(k)}-(1-\lambda_k)y^{(k)})\rangle\\ &= \lambda_k\langle\nabla g(p^{(k)}), y^*-p^{(k)}\rangle+(1-\lambda_k)\langle\nabla g(p^{(k)}), y^{(k)}-p^{(k)}\rangle\\ &\le \lambda_k\left[g(y^*)-g(p^{(k)})\right]+(1-\lambda_k)\left[g(y^{(k)})-g(p^{(k)})\right].\end{aligned}$$

We are now ready to analyze the decrease of the objective function in each iteration of Algorithm 2. The expectation in the following equations is performed with respect to $C_{i_k}\sim P$.

$$\begin{aligned}&\mathbb{E}\left[g(y^{(k+1)})\right]\\ &\overset{1)}{\le} g(p^{(k)})+\frac{\lambda_k}{\alpha}\mathbb{E}\left\{\sum_{r\in C_{i_k}}\left[\langle\nabla_r g(p^{(k)}), z_r^{(k+1)}-z_r^{(k)}\rangle+\frac{\lambda_k}{2\alpha}\|z_r^{(k+1)}-z_r^{(k)}\|_{2,\mu^C}^2\right]\right\}\\ &\overset{2)}{=} g(p^{(k)})+\frac{\lambda_k}{\alpha}\mathbb{E}\left\{\sum_{r\in C_{i_k}}\left[\langle\nabla_r g(p^{(k)}), z_r^{(k+1)}-z_r^{(k)}\rangle+\frac{\lambda_k}{2\alpha}\|z_r^{(k+1)}-z_r^{(k)}\|_{2,\theta_r^P}^2\right]\right\}\\ &= g(p^{(k)})+\frac{\lambda_k}{\alpha}\mathbb{E}\left\{\sum_{r\in C_{i_k}}\left[\langle\nabla_r g(p^{(k)}), y_r^*-z_r^{(k)}\rangle+\langle\nabla_r g(p^{(k)}), z_r^{(k+1)}-z_r^*\rangle+\frac{\lambda_k}{2\alpha}\|z_r^{(k+1)}-z_r^{(k)}\|_{2,\theta_r^P}^2\right]\right\}\\ &\overset{3)}{\le} g(p^{(k)})+\frac{\lambda_k}{\alpha}\mathbb{E}\left\{\sum_{r\in C_{i_k}}\left[\langle\nabla_r g(p^{(k)}), y_r^*-z_r^{(k)}\rangle-\frac{\lambda_k}{2\alpha}\|z_r^{(k+1)}-y_r^*\|_{2,\theta_r^P}^2+\frac{\lambda_k}{2\alpha}\|y_r^*-z_r^{(k)}\|_{2,\theta_r^P}^2\right]\right\}\\ &= g(p^{(k)})+\lambda_k\langle\nabla g(p^{(k)}), y^*-z^{(k)}\rangle+\frac{\lambda_k^2}{2\alpha^2}\mathbb{E}\left[\|z_{[C_{i_k}]}^{(k)}-y_{[C_{i_k}]}^*\|_{2,\theta^P}^2-\|z_{[C_{i_k}]}^{(k+1)}-y_{[C_{i_k}]}^*\|_{2,\theta^P}^2\right]\\ &\overset{4)}{=} g(p^{(k)})+\lambda_k\langle\nabla g(p^{(k)}), y^*-z^{(k)}\rangle+\frac{\lambda_k^2}{2\alpha^2}\mathbb{E}\left[\|z^{(k)}-y^*\|_{2,\theta^P}^2-\|z^{(k+1)}-y^*\|_{2,\theta^P}^2\right]\\ &\overset{5)}{=} g(y^*)+(1-\lambda_k)\left[g(y^{(k)})-g(y^*)\right]+\frac{\lambda_k^2}{2\alpha^2}\left\{\|z^{(k)}-y^*\|_{2,\theta^P}^2-\mathbb{E}\left[\|z^{(k+1)}-y^*\|_{2,\theta^P}^2\right]\right\},\end{aligned}\tag{17}$$

where 1) follows from (8), 2) may be deduced from Lemma G.1, 3) is a consequence of Lemma J.2 and Lemma J.3, 4) is due to the fact that $y_r^{(k+1)}=y_r^{(k)}$ for $r\notin C_{i_k}$, and 5) follows from Lemma J.4.

Based on the definition of $\{\lambda_k\}_{k\ge0}$, we also have

$$\frac{1-\lambda_k}{\lambda_k^2}=\frac{1}{\lambda_{k-1}^2},\quad 0<\lambda_{k+1}\le\lambda_k\le\frac{2}{k+2/\lambda_0}=\frac{2}{k+2}.\tag{18}$$

Hence, combining the above expression with (17), for $k\in[1,\frac{2}{\alpha}\lceil\sqrt{N\|\theta^P\|_{1,\infty}}\rceil+1]$, we have

$$\begin{aligned}&\mathbb{E}\left[\frac{1-\lambda_k}{\lambda_k^2}\left[g(y^{(k)})-g(y^*)\right]+\frac{1}{2\alpha^2}\|z^{(k)}-y^*\|_{2,\theta^P}^2\right]\\ &=\mathbb{E}\left[\frac{1}{\lambda_{k-1}^2}\left[g(y^{(k)})-g(y^*)\right]+\frac{1}{2\alpha^2}\|z^{(k)}-y^*\|_{2,\theta^P}^2\right]\\ &\le\mathbb{E}\left[\frac{1-\lambda_{k-1}}{\lambda_{k-1}^2}\left[g(y^{(k-1)})-g(y^*)\right]+\frac{1}{2\alpha^2}\|z^{(k-1)}-y^*\|_{2,\theta^P}^2\right]\\ &\le\cdots\le\frac{(1-\lambda_0)}{\lambda_0^2}\left[g(y^{(0)})-g(y^*)\right]+\frac{1}{2\alpha^2}\|z^{(0)}-y^*\|_{2,\theta^P}^2.\end{aligned}\tag{19}$$

Lemma 3.9 implies the strong convexity property as

$$\begin{aligned}
\|Ay^{(k)} - Ay^*\|_2^2 &\geq \frac{2}{N\|\theta^P\|_{1,\infty}}\|y^{(k)} - y^*\|_{2,\theta^P}^2 \Rightarrow \\
g(y^{(k)}) - g(y^*) &\geq \langle \nabla g(y^*), y^{(k)} - y^* \rangle + \frac{1}{N\|\theta^P\|_{1,\infty}}\|y^{(k)} - y^*\|_{2,\theta^P}^2 \\
&\overset{1)}{\geq} \frac{1}{N\|\theta^P\|_{1,\infty}}\|y^{(k)} - y^*\|_{2,\theta^P}^2. \qquad (20)
\end{aligned}$$

Here, 1) holds since y^* is an optimal solution of $\min_y g(y)$ and thus $\langle \nabla g(y^*), y^{(k)} - y^* \rangle \geq 0$.

Combining (18), (19) and (20), we obtain

$$\begin{aligned}
\mathbb{E}\left[g(y^{(k)}) - g(y^*)\right] &\leq \lambda_{k-1}^2 \left[\frac{1-\lambda_0}{\lambda_0^2}(g(y^{(0)}) - g(y^*)) + \frac{1}{2\alpha^2}\|y^{(0)} - y^*\|_{2,\theta^P}^2\right] \\
&\leq \left(\frac{2}{k+1}\right)^2 \frac{1}{2\alpha^2}\|y^{(0)} - y^*\|_{2,\theta^P}^2 \\
&\leq \left(\frac{2}{k+1}\right)^2 \frac{N\|\theta^P\|_{1,\infty}}{2\alpha^2}(g(y^{(0)}) - g(y^*)).
\end{aligned}$$

Therefore, for $k = \left\lceil (1+c)\frac{\sqrt{2N\|\theta^P\|_{1,\infty}}}{\alpha} + c \right\rceil$, we have

$$\mathbb{E}\left[g(y^{\left(\left\lceil (1+c)\frac{\sqrt{2N\|\theta^P\|_{1,\infty}}}{\alpha}+c\right\rceil\right)}) - g(y^*)\right] \leq \frac{1}{1+c}(g(y^{(0)}) - g(y^*)).$$

For each value of $k = l \times \left\lceil (1+c)\frac{\sqrt{2N\|\theta^P\|_{1,\infty}}}{\alpha} + c \right\rceil$, $l \in \mathbb{Z}_{\geq 0}$, the values $z^{(k)}$, λ_k are reinitialized. Using a similar proof as above, we have

$$\mathbb{E}\left[g(y^{\left((l+1)\times\left\lceil (1+c)\frac{\sqrt{2N\|\theta^P\|_{1,\infty}}}{\alpha}+c\right\rceil\right)}) - g(y^*)\right] \leq \frac{1}{1+c}\left[g(y^{\left(l\times\left\lceil (1+c)\frac{\sqrt{2N\|\theta^P\|_{1,\infty}}}{\alpha}+c\right\rceil\right)}) - g(y^*)\right].$$

Therefore,

$$\mathbb{E}\left[g(y^{\left(l\left\lceil (1+c)\frac{\sqrt{2N\|\theta^P\|_{1,\infty}}}{\alpha}+c\right\rceil\right)}) - g(y^*)\right] \leq \frac{1}{(1+c)^l}(g(y^{(0)}) - g(y^*)).$$

This concludes the proof.

J.2 Avoiding Full-Dimensional Vector Operations

Algorithm 2 can be implemented without full-dimensional vector operations. In each step, only those coordinates within the blocks in C are updated. Consequently, one only needs to replace $p^{(k)}$ and $y^{(k)}$ with $p^{(k)} = z^{(k)} + \lambda_k^2 u^{(k)}$ and $y^{(k)} = z^{(k)} + \lambda_{k-1}^2 u^{(k)}$, where $u^{(k)}$ is a new variable described in Algorithm 3.

K Minimization of $\|\theta^P\|_{1,\infty}$

We first define $\triangle_* \triangleq \max_{r\in[R]} |\{r' \in [R] | S_{r'} \cap S_r \neq \emptyset\}|$, which we use in our subsequent derivations.

As shown in Theorem 3.10 and Theorem J.1, $\|\theta^P\|_{1,\infty}$ plays an important role in the convergence rate of CDMs. Hence, we are interested in identifying the optimal sampling strategy P that minimizes $\|\theta^P\|_{1,\infty}$.

Algorithm 3: Parallel ACDM for Solving Problem (7) (an efficient implementation)

Input: $\mathcal{B}$, α

0: Initialize $z^{(0)} \in \mathcal{B}$, $u^{(0)} \leftarrow 0 \in \mathbb{R}^N$, $k \leftarrow 0$.

1: Do the following steps iteratively until the dual gap $< \epsilon$:

2: If $k = l \left\lceil (1+c)\sqrt{\frac{2N\|\theta^P\|_{1,\infty}}{\alpha}} + c \right\rceil$ for some $l \in \mathbb{Z}$ and $c > 0$,
$z^{(k)} \leftarrow z^{(k)} + \lambda_{k-1}^2 u^{(k)}$, $u^{(k)} \leftarrow 0$, $\lambda_k \leftarrow 1$

3: Sample one set C_{i_k} according to a α-proper distribution P

4: For $r \in C_{i_k}$:

5: $\triangle z_r \leftarrow \arg\min_{\triangle z + z_r^{(k)} \in B_r} \|\triangle z + \frac{\alpha}{\lambda_k}(\theta_r^P)^{-1} \odot \nabla_r g(z^{(k)} + \lambda_k^2 u^{(k)})\|_{2,\theta_r^P}^2$

6: $z_r^{(k+1)} \leftarrow z_r^{(k)} + \triangle z_r$

7: $u_r^{(k+1)} \leftarrow u_r^{(k)} + \frac{\lambda_k - \alpha}{\alpha \lambda_k^2} \triangle z_r$

8: $\lambda_{k+1} \leftarrow \frac{\sqrt{\lambda_k^4 + 4\lambda_k^2} - \lambda_k^2}{2}$

9: $k \leftarrow k + 1$

10: Output $y^{(k)}$

Algorithm 4: A Greedy Algorithm for Finding a Balanced-Partition Distribution

Input: $\{S_r\}_{r\in[R]}$, K

0: Initialize the partition $\mathcal{C} = \{C_i\}_{1\le i\le m}$, $C_i \leftarrow \emptyset$, vectors $\{\mu^{C_i}\}_{1\le i\le m}$, $\mu^{C_i} \in \mathbb{R}^N$, and $\mu^{\max} \in \mathbb{R}^N$, $\mu^{\max} \leftarrow 0$.

1: For r from 1 to R:

2: For i from 1 to m:

3: If $|C_i| < K$:

4: $\triangle\mu^{C_i} \leftarrow 0$

5: For v in S_r, if $\mu_v^{C_i}$ is equal to $\mu_v^{\max}$, $\triangle\mu^{C_i} \leftarrow \triangle\mu^{C_i} + 1$

6: else: $\triangle\mu^{C_i} \leftarrow \infty$

7: $i^* \leftarrow \arg\min_i \triangle\mu^{C_i}$

8: $C_{i^*} \leftarrow C_{i^*} \cup \{r\}$

9: For v in S_r, $\mu_v^{C_{i^*}} \leftarrow \mu_v^{C_{i^*}} + 1$, $\mu_v^{\max} \leftarrow \max\{\mu_v^{\max}, \mu_v^{C_{i^*}}\}$.

10: Output $\mathcal{C}$.

For this purpose, consider a partition of $[R]$ into $m = \lceil \frac{1}{\alpha} \rceil$ parts $\{C_i\}_{1\le i\le m}$, such that $|C_i| \in \{K-1, K\}$. We refer to such a partition as a *balanced partition*. In this case, every block r is in exactly one component C_i and $\|\theta^P\|_{1,\infty} = \sum_{v\in[N]} \max_{i\in[m]} \mu_v^{C_i}$. As a result, the problem of minimizing $\|\theta^P\|_{1,\infty}$ is closely related to the so called *equitable coloring problem* first proposed by Meyer [20].

Definition K.1 (Meyer [20])**.** Given a graph, an *equitable coloring* is an assignment of colors to the vertices that satisfies the following two properties: no two adjacent vertices share the same color and the number of vertices in any two color classes differs by at most one. Moreover, the minimum number of colors in any equitable coloring is termed the *equitable coloring number*.

Hajnal-Szemerédi's Theorem [25] established one of the most important results in equitable graph coloring: a graph is equitably k-colorable if k is strictly greater than the maximum vertex degree. This bound is tight. We can construct a graph based on the incidence structure of DSFM problem so that a vertex corresponds to a component submodular function and two vertices are connected iff the corresponding submodular functions are incident to at least one common point. An equitable coloring of this graph can be used to assign submodular functions of the same color class to a set C_i in $\mathcal{C}$. This guarantees that $\mu_v^{C_i} \le 1$ for all C_i and all $v \in [N]$. Note that the maximal degree of this graph is $\triangle_*$. By directly applying Hajnal-Szemerédi's Theorem, we have the following lemma.

Lemma K.2. There exists a balanced-partition distribution P such that $\|\theta^P\|_{1,\infty} = N$, provided that $\lceil \frac{1}{\alpha} \rceil \ge \triangle_* + 1$.

As in many applications, such as image segmentation [9], the value of $\triangle_*$ is small, and hence using a balanced-partition instead of one obtained through sampling uniformly at random may produce significantly better results. Unfortunately, finding the equitable coloring number is an NP-hard

problem; still, a polynomial time algorithm for finding $\triangle_* + 1$ equitable colorings was described in [31], with complexity $O(\triangle_* R^2)$. We describe a greedy algorithm that outputs a balanced-partition distribution and aims to minimize $\|\theta^P\|_{1,\infty}$ in Algorithm 4. According to our experimental results, the sampling strategy P found by Algorithm 4 works better than sampling uniformly at random.

L Using Weighted Proximal Terms

The AP and RCDM solvers discussed in the main text are designed to solve the convex optimization (2), but also produce a solution to the discrete optimization problem (1). To solve the discrete optimization problem (1), another convex optimization formulation may be considered instead:

$$\min_{x\in\mathbb{R}^N} \sum_{r\in[R]} f_r(x) + \frac{1}{2}\|x\|^2_{2,w}, \tag{21}$$

where the choice of $w \in \mathbb{R}^N_{>0}$ will be described later. By using the arguments in [32] or in Chapter 8.1-8.2 of [12], we know that the solution of the discrete optimization problem (1) can be obtained as $S = \{i \in [N] | x^*_i > 0\}$, where x^* is a solution of (21).

Next, we describe how a proper choice of w allows one to avoid compute oblique projections in the AP and parallel CDM algorithms. If oblique projections are allowed, a good choice for w may also decrease the computational complexities listed in Table 1. The results obtained based on weighted proximal terms are summarized in Table 2.

	Using Orthogonal Projection $\Pi_{\mathcal{B}_r}(\cdot)$	
	The Value of w	Complexity
AP	$w = \mu$	$O(N\|\mu\|_1 \frac{R}{K})$
RCDM	$w = \frac{R-K}{R-1}1 + \frac{K-1}{R-1}\mu$	$O\left(\left(\frac{R-K}{R-1}N^2 + \frac{K-1}{R-1}N\|\mu\|_1\right)\frac{R}{K}\right)$
ACDM	$w = \frac{R-K}{R-1}1 + \frac{K-1}{R-1}\mu$	$O\left(\left(\frac{R-K}{R-1}N^2 + \frac{K-1}{R-1}N\|\mu\|_1\right)^{\frac{1}{2}}\frac{R}{K}\right)$
	Using Oblique Projection $\Pi_{\mathcal{B}_r, w^{1/2}}(\cdot)$	
	The Value of w	Complexity
AP	$w = \mu^{\frac{1}{2}}$	$O(\|\mu^{\frac{1}{2}}\|_1^2 \frac{R}{K})$
RCDM	$w = \left(\frac{R-K}{R-1}1 + \frac{K-1}{R-1}\mu\right)^{\frac{1}{2}}$	$O\left(\left\|\left(\frac{R-K}{R-1}1 + \frac{K-1}{R-1}\mu\right)^{\frac{1}{2}}\right\|_1^2 \frac{R}{K}\right)$
ACDM	$w = \left(\frac{R-K}{R-1}1 + \frac{K-1}{R-1}\mu\right)^{\frac{1}{2}}$	$O\left(\left\|\left(\frac{R-K}{R-1}1 + \frac{K-1}{R-1}\mu\right)^{\frac{1}{2}}\right\|_1 \frac{R}{K}\right)$

Table 2: New complexity results based on weighted proximal terms: here, complexity refers to the required number of iterations needed to achieve an $\epsilon-$optimal solution (the dependence on ϵ is the same for all algorithms and hence omitted). As before, K is the parallelization parameter and it equals the number of min-norm points problems that are solved within each iteration; $K = 1$ reduces to the sequential case.

We now analyze the new objective (21) in more detail. The proof techniques used in the main text carry over to the setting involving weighted proximal terms.

By using a dual strategy similar to those described in Lemma 2.1 and Lemma 3.2, we arrive at the dual formulation of problem (21) described in the next lemma. Note that the derivation of (L.1) takes into account the underlying incidence relations.

Lemma L.1. The dual problem of (21) reads as

$$\min_{a,y} \|a - y\|^2_{2,I(w^{-1}\odot\mu)} \quad \text{s.t.} \quad y \in \mathcal{B}, Aa = 0, \text{ and } a_{r,i} = 0,\ \forall (r,i) : i \notin S_r, r \in [R]. \tag{22}$$

Moreover, problem may be written in a more compact form as

$$\min_{y} \|Ay\|^2_{2,w^{-1}} \quad \text{s.t.} \quad y \in \mathcal{B}. \tag{23}$$

For both problems, the primal and dual variables are related according to $x = -w^{-1} \odot Ay$.

L.1 The Incidence Relations AP (IAP) Method for Solving (L.1)

The steps of the IAP method are listed in Algorithm 5.

Algorithm 5: The IAP Method for Solving (L.1)

0: For all r, initialize $y_r^{(0)} \in \mathcal{B}_r$, and $k \leftarrow 0$
1: In iteration k:
2: For all $r \in [R]$:
3: $a_{r,i}^{(k+1)} \leftarrow y_{r,i}^{(k)} - \mu_i^{-1}(Ay^{(k)})_i$ for all $i \in S_r$
4: $y_r^{(k+1)} \leftarrow \Pi_{\mathcal{B}_r, w^{-1}\odot\mu}(a_r^{(k+1)})$

The convergence properties of Algorithm 5 can be characterized similarly as those of IAP for solving (5). The latter relies on a finite upper bound for $\kappa_* \triangleq \sup_{y\in\mathcal{Z}\cup\mathcal{B}/\Xi} \frac{d_{I(w^{-1}\odot\mu)}(y,\Xi)}{\max\{d_{I(w^{-1}\odot\mu)}(y,\mathcal{Z}), d_{I(w^{-1}\odot\mu)}(y,\mathcal{B})\}}$.

Lemma L.2. One has $\kappa_* \leq \sqrt{\frac{\|w^{-1}\odot\mu\|_1\|w\|_1}{2}} + 1$. When $w = \mu$, $\kappa_* \leq \sqrt{\frac{N\|\mu\|_1}{2}} + 1$.

Proof. The result follows using the same strategy as the one used to prove Lemma 3.4. Note that when using Lemma 3.1, one should set θ to $I(w^{-1} \odot \mu)$ and replace w by w^{-1}. □

By setting $w = \mu$, Step 4 of Algorithm 5 reduces to orthogonal projections. In this case, based on Lemma L.2, Algorithm 5 requires $O(N\|\mu\|_1 \log \frac{1}{\epsilon})$ iterations to achieve an ϵ-optimal solution. By setting $w = \mu^{\frac{1}{2}}$ for all $i \in [R]$, Step 4 of Algorithm 5 reduces to the projections $\Pi_{\mathcal{B}_r, w^{\frac{1}{2}}}(\cdot)$. In this case, Algorithm 5 requires $O\left(\left\|\mu^{\frac{1}{2}}\right\|^2 \log \frac{1}{\epsilon}\right)$ iterations to achieve an ϵ-optimal solution. The latter result is slightly better because $\left\|\mu^{\frac{1}{2}}\right\|^2 \leq N\|\mu\|_1$.

L.2 A Parallel RCD Method for Solving (23) with Uniform Sampling Strategies

As discussed in Section 3.3, RCDM strongly depends on an α-proper distribution P that characterizes the parallel coordinate sampling strategy. In what follows, we choose P to be a uniform distribution. From Lemma 3.12, we know that when P is uniform, one has $\theta_r^P = \frac{K-1}{R-1}\mu + \frac{R-K}{R-1}1$ for all $r \in [R]$, where K denotes the number of projections computed in parallel as part of each iteration. In Algorithm 1, θ_r^P defines the normed space over which to minimize $g(y)$. As our goal is to minimize $g_w(y) = \frac{1}{2}\|Ay\|_{2,w^{-1}}^2$, the vector used to define the normed space is

$$\nu = w^{-1} \odot \theta_r^P = w^{-1} \odot (\frac{K-1}{R-1}\mu + \frac{R-K}{R-1}1).$$

The parallel RCDM procedure in this setting is described in Algorithm 6.

Algorithm 6: Parallel RCDM with Uniform Sampling for Solving (23)

Input: $\mathcal{B}$, K
0: Initialize $y^{(0)} \in \mathcal{B}$, $k \leftarrow 0$
1: Do the following steps iteratively until the dual gap $< \epsilon$:
2: Uniformly sample $C_{i_k} \subseteq [R]$ so that $|C_{i_k}| = K$.
3: For $r \in C_{i_k}$:
4: $y_r^{(k+1)} \leftarrow \Pi_{\mathcal{B}_r,\nu}(y_r^{(k)} - (\nu^{-1}) \odot \nabla_r g_w(y))$
5: Set $y_r^{(k+1)} \leftarrow y_r^{(k)}$ for $r \notin C_{i_k}$
6: $k \leftarrow k + 1$
7: Output $y^{(k)}$

Similarly to what was done in Lemma 3.9, we can establish weak strong convexity of $g_w(y)$ with respect to the norm $\|\cdot\|_{2,\nu}$ by invoking Lemma 3.1.

Lemma L.3. For any $y \in \mathcal{B}$, let $y^* = \arg\min_{\xi \in \Xi} \|\xi - y\|_{2,\nu}^2$. Then,

$$\|Ay - Ay^*\|_{2,w^{-1}}^2 \geq \frac{2}{\|w\|_1 \|\nu\|_1} \|y - y^*\|_{2,\nu}^2.$$

Therefore, using a strategy similar to the one outlined in the proof of Theorem 3.10, the convergence rates of Algorithm 6 can be derived as summarized in the next theorem.

Theorem L.4. At each iteration of Algorithm 6, $y^{(k)}$ satisfies

$$\mathbb{E}\left[g_w(y^{(k)}) - g_w(y^*) + \frac{1}{2} d_{I(\nu)}^2(y^k, \xi)\right] \leq \left[1 - \frac{4K}{R(\|w\|_1 \|\nu\|_1 + 2)}\right]^k \left[g_w(y^{(0)}) - g_w(y^*) + \frac{1}{2} d_{I(\nu)}^2(y^0, \xi)\right].$$

By setting $w = \frac{K-1}{R-1}\mu + \frac{R-K}{R-1}1$, we reduce the projections in Step 4 of Algorithm 6 to orthogonal projections. In this case, based on Theorem L.4, Algorithm 6 requires $O\left(\left(\frac{K-1}{R-1} N\|\mu\|_1 + \frac{R-K}{R-1} N^2\right) \frac{R}{K} \log \frac{1}{\epsilon}\right)$ iterations to achieve an ϵ-optimal solution.

By setting $w = \left(\frac{K-1}{R-1}\mu + \frac{R-K}{R-1}\right)^{1/2}$ for all $i \in [R]$, the projections in Step 4 of Algorithm 6 reduce to oblique projections $\Pi_{\mathcal{B}_r, w^{\frac{1}{2}}}(\cdot)$. In this case, Algorithm 6 requires $O\left(\left\|\left(\frac{K-1}{R-1}\mu_i + \frac{R-K}{R-1}\right)^{1/2}\right\|_1^2 \log \frac{1}{\epsilon}\right)$ iterations to achieve an ϵ-optimal solution, which is slightly better than the previous case. The accelerated methods can be analyzed in the same manner.

L.3 Simulations

We now describe simulation results that empirically evaluate Algorithms 5 and 6. The DSFM problem is designed as follows. We consider $N = 100$ vertices. The unary potentials of different elements are iid standard Gaussian variables. We construct a network over these vertices based on the Barabási-Albert model (BA) [33], initialized with a single edge between vertices 1 and 2. Each edge in the network gives a pairwise potential for the corresponding vertices. We use the BA model so that the number of incidence relations corresponding to different vertices vary to a large extent. As we are using weighted proximal terms, the continuous objectives are not consistent for different w. However, here, we are only interested in generating solutions for the discrete problem (1) and thus regard the discrete gap ν_d as the relevant metric for characterizing convergence properties. The following results are obtained from 100 independent experiments.

In IAP (Algorithm 5), we set $w \in \{1, \mu, \mu^{1/2}\}$, corresponding to three cases: *unweighted proximal term + oblique projections*, *weighted proximal term + orthogonal projections*, *weighted proximal term + oblique projections*, respectively. In RCDM-U (Algorithm 6), we set $w \in \{1, \frac{K-1}{R-1}\mu + \frac{R-K}{R-1}1, (\frac{K-1}{R-1}\mu + \frac{R-K}{R-1}1)^{1/2}\}$, corresponding to the same three cases. We control the number of parallel projection operations in each iteration by choosing $K \in \{10, 20, 30, 40, 50\}$. Figure 5 shows the convergence curve of the discrete gap for different solvers and different choices of w. We only plotted results for $K = 10, 50$ as other values of K produce similar patterns. For both IAP and RCDM-U, when w corresponds to the weighted proximal term + orthogonal projections case, we obtain the best convergence rates. The value $w = 1$, corresponding to the case unweighted proximal term + oblique projections, results in the worst convergence rates. Albeit somewhat inconsistent with the results listed in Table 2, the simulations simply imply that using weighted proximal terms can reduce the complexity of the algorithms at hand and that the weighted proximal term with orthogonal projections in the inner loop may represent the best choice in practice.

In Table 3, we also list the number of iterations needed by different solvers to obtain a solution for the discrete problem (1). Again, the w corresponding to the weighted proximal term + orthogonal projections case results in the smallest number of iterations, while the w corresponding to the unweighted proximal term + oblique projections case results in the largest number of iterations. Note that as K increases, the number of iterations $\times K/R$ in IAP does not change as IAP is fully parallelizable, while the number of operations in RCDM-U slightly increases due to the overlapping incidence sets of different submodular functions.

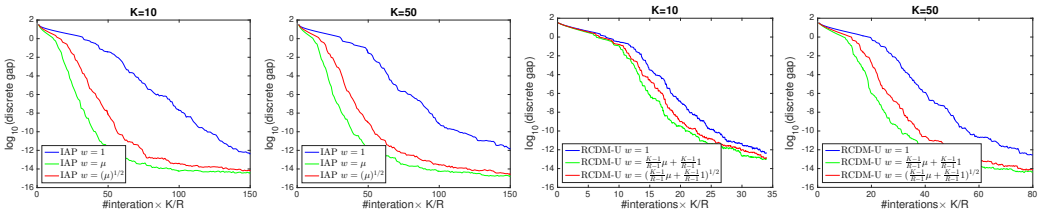

Figure 5: Simulations for Algorithm 5 and 6: $\log_{10}(\text{discrete gap})$ vs (number of iterations $\times K/R$).

Solvers	w	$K=10$		$K=20$		$K=30$		$K=40$		$K=50$	
		MN	MD	MN	MD	MN	MD	MN	MD	MN	MD
IAP	1	109	103	109	103	109	103	109	103	109	103
	μ	43	34	43	34	43	34	43	34	43	34
	$\mu^{1/2}$	59	50	59	50	59	50	59	50	59	50
RCDM-U	1	27	22	34	28	43	38	51	46	54	49
	$\frac{K-1}{R-1}\mu + \frac{R-K}{R-1}1$	22	17	25	20	29	24	32	24	33	25
	$\left(\frac{K-1}{R-1}\mu + \frac{R-K}{R-1}1\right)^{1/2}$	25	19	28	23	33	28	37	31	38	32

Table 3: The number of iterations $\times K/R$ needed to find an optimal solution to the discrete problem (1). MN: mean; MD: median.

M Supplementary experiments

Semi-supervised learning over hypergraphs. We also evaluate the proposed approaches over the 20Newsgroups from the University of California Irvine (UCI) data repository. This dataset is used as a benchmark example for evaluating semisupervised learning algorithms over hypergraphs [34, 35]. Here, for simplicity, we focused on binary classification tasks and thus paired the four 20Newsgroups classes, so that one group includes "Comp." and "Sci", and the other one includes "Rec." and "Talk". The 20Newsgroups dataset consists of categorical features and we adopt the same approach as the one described in [34] to construct hyperedges: each feature corresponds to one hyperedge and contributes one submodular function to the decomposition. Hence, 20Newsgroups contains $N = 16242$ elements and $R = 100$ submodular functions.

In the experiments for 20Newsgroups, we uniformly at random picked 200 elements and set their corresponding components in x_0 of equation (9) to the true labels and set all other entries to zero. Figure 6 shows the results of the experiments pertaining to 20Newsgroup. We compared the convergence rate of different algorithms for different values of the parameter $\alpha \in \{0.02, 0.1\}$. The value on the horizontal axis, # iterations $\times\alpha$, equals the total number of projections, scaled by R. The results are averaged over 10 independent experiments. Once again, we observe that CD-based methods outperform AP-based methods. ACDM-U offers the best performance among all CD-based methods and IAP significantly outperforms AP. Similarly, RCDM-G has better performance than RCDM-U, due to the use of the greedy algorithm for the sampling procedure.

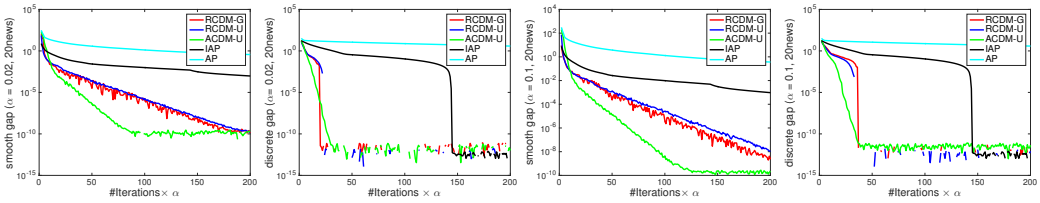

Figure 6: 20Newsgrounp: Smooth/discrete gap vs the (number of iterations $\times\alpha$).