Constrained Robust Submodular Partitioning

Shengjie Wang[*,1],Tianyi Zhou[*,1,2], Chandrashekhar Lavania[1] & Jeff A. Bilmes[1]

University of Washington, Seattle[1]; University of Maryland, College Park[2]
{wangsj,tianyizh,lavaniac,bilmes}@uw.edu

Abstract

In the robust submodular partitioning problem, we aim to allocate a set of items into m blocks, so that the evaluation of the minimum block according to a submodular function is maximized. Robust submodular partitioning promotes the diversity of every block in the partition. It has many applications in machine learning, e.g., partitioning data for distributed training so that the gradients computed on every block are consistent. We study an extension of the robust submodular partition problem with additional constraints (e.g., cardinality, multiple matroids, and/or knapsack) on every block. For example, when partitioning data for distributed training, we can add a constraint that the number of samples of each class is the same in each partition block, ensuring data balance. We present two classes of algorithms, i.e., Min-Block Greedy based algorithms (with an $\Omega(1/m)$ bound), and Round-Robin Greedy based algorithms (with a constant bound) and show that under various constraints, they still have good approximation guarantees. Interestingly, while normally the latter runs in only weakly polynomial time, we show that using the two together yields strongly polynomial running time while preserving the approximation guarantee. Lastly, we apply the algorithms on a real-world machine learning data partitioning problem showing good results.

1 Introduction

The problem of partitioning a given set V of items into m blocks, where any two blocks share no items in common, arises in many real-world scenarios and machine learning applications. As an optimization problem, partitioning aims to generate the blocks so that the utilities of the blocks, as measured by a given set function, are good. Submodular functions are a rich family of set functions that naturally captures diversity of a given set of items. They have been applied in many real-world problems [26, 30, 21, 17, 13, 18, 28]. By maximizing a submodular utility function for each partitioned block, we encourage each block to be representative of the ground set V. Many algorithms have been proposed for various submodular partitioning problems with approximation guarantees.

For the submodular welfare problem [24], we aim to find a partition such that the sum of the submodular evaluations of every block is maximized. Such an objective promotes the overall utility of the entire partition but some blocks may still have small function values. The robust submodular partitioning problem [12, 29] (often called "submodular fair allocation with indivisible goods") aims to find the partition such that the minimum-valued block in the partition is maximized according to the submodular function. The robust objective optimizes the worst block in the partition so that all blocks are minimally "good." In the general setting, every block in the partition may have a different submodular function (the heterogeneous case) although for this work, we study only the restricted setting where all blocks share the same submodular function (the homogeneous case). The robust submodular partitioning problem has many applications. Given V as the training dataset for a machine learning task, Wei et al. [29] finds a partition of V for distributed training: every block of partitioned data is sent to a single machine for parallel gradient computations, and the gradients

35th Conference on Neural Information Processing Systems (NeurIPS 2021), Sydney, Australia.

are aggregated over all the blocks in the partition for model updates. Since we enforce each block to be representative of V, the gradients computed across distributed machines are consistent, resulting in reduced variance and improved convergence for the aggregation step. Using a similar idea, Wang et al. [25] partitions the training data into mini-batches so that every mini-batch is as representative as possible, therefore reducing the variance during mini-batch gradient-based training.

In this work, we explore two different algorithmic approaches, *Min-Block Greedy* and *Round-Robin Greedy*, for our partitioning problem but under various constraints, newly applied to this problem. For Min-Block Greedy based algorithms, we first show that the $\frac{1}{m}$ bound for the unconstrained case is tight. We then modify the algorithm to allow a general down-closed constraint $\mathcal{C}$, and prove an approximation bound of $\frac{\alpha}{\alpha m+1}$, where α is the bound for solving the submodular maximization problem under constraint $\mathcal{C}$ using a greedy based algorithm. For example, for a cardinality constraint, $\alpha = 1-1/e$ [9], and the bound for constrained robust submodular partitioning is $(m + \frac{1}{1-1/e})^{-1}$. Similarly, for $\mathcal{C}$ as an intersection-of-p-matroids constraint, $\alpha = (1+p)^{-1}$ [10], and the bound is $(m+p+1)^{-1}$; for $\mathcal{C}$ as a knapsack constraint, $\alpha = 0.5(1-1/e)$ [16], and the bound is $(m + \frac{2}{1-1/e})^{-1}$. For Round-Robin Greedy based algorithms, when $\mathcal{C}$ is a cardinality constraint, we get a bound of $\frac{(1-1/e)^2}{3}$, and when $\mathcal{C}$ is a matroid constraint, we get a bound of $\frac{1-1/e}{5}$. The Min-Block Greedy approach gives a weaker bound, and since the $\frac{1}{m}$ bound for the unconstrained case is tight, we cannot improve upon the $\frac{1}{m}$ factor for the constrained case. The Round-Robin Greedy approach gives a constant bound, but its running time is worse. The running time for Min-Block Greedy is $\mathcal{O}(n^2)$, where n is the ground set size. For Round-Robin Greedy under a matroid constraint, the running time is $\mathcal{O}(n^2(\log\log m + \log\frac{1}{\delta}))$, as it needs to use binary search to find the optimal solution value to the given problem over an exponentially decreasing sequence, with $\frac{1}{1+\delta}$ ($\delta > 0$) as the multiplicative factor. In all cases, we assume an oracle model, and the running time is in terms of the number of submodular evaluations. An important contribution our work shows is that by utilizing the Min-Block Greedy algorithm result as input, our Round-Robin Greedy algorithm attains **strongly polynomial running time** — all previous results on the unconstrained case using a Round-Robin-like algorithm have only weakly polynomial running time [3].

The various constraints (e.g., cardinality, matroids, and knapsack) we study greatly improves the applicability of robust submodular partitioning. Several applications that benefit from the constraints include: (1) Partition a training data for machine learning models in distributed training or forming deterministic mini-batches [29, 25]. The additional constraint can be the number of samples from each class to be no more than a certain value. If there are enough samples in the training data, every resulting block will have the same number of samples for each class, which avoids imbalance, further promoting each block's diversity, and improving the gradients' consistency. (2) Given an undirected graph, we partition the edges into subgraphs so that each subgraph is representative based on the submodular evaluation, and we also constrain each subgraph to have no cycles (a cycle matroid). A practical scenario is that we wish to send information efficiently over a graph of devices. We partition the graph so that information can be sent in parallel, and the constraint to have no cycles enforces that information is not redundantly sent twice to the same device, leading to improved communications efficiency. (3) Again, for an undirected and connected graph, we partition the edges into subgraph blocks such that if we were to remove any block of the partition from the original graph, the remaining graph remains connected (which can be done via a bond matroid, where min-cuts are cycles and anything not a cut is independent). In practice, this functions as a form of reliability insurance. For a graph of devices, we partition the graph to perform computation in parallel, so that if the connections in one block fail, the other blocks can still operate and communicate since the graph remains connected.

2 Related Work

Golovin [12] introduces robust submodular partitioning (i.e., submodular fair allocation of indivisible goods), and proposes a matching-based algorithm with a bound of $\frac{1}{n-m+1}$. Khot & Ponnuswami [14] proposes a binary search based algorithm and gives an improved bound of $\frac{1}{2m-1}$. Asadpour & Saberi [2] uses an ellipsoid approximation approach and gives a bound of $\Omega(\frac{1}{\sqrt{n}m^{1/4}\log n \log^{3/2} m})$. Wei et al. [29] gives a simple Min-Block Greedy algorithm and proves a $\frac{1}{m}$ bound. A Round-Robin Greedy approach is given in [3] with a bound of $\frac{1-e^{-1}}{3}$. Ghodsi et al. [11] proposes a local search algorithm with a bound of $\frac{1}{3}$. Both [3] and [11] requires guessing of the optimal solution value from an exponentially decreasing sequence of values, so strictly speaking, they lose an extra $(1+\delta)$-factor in the approxima-

tion bound where $(1+\delta)$ is the exponential factor for the guessing sequence. We can set the δ value small to get close to the constant bounds shown above at the costs of computation. Wang et al. [25] extends the Min-Block Greedy algorithm with a cardinality constraint, and also shows a hierarchical partitioning framework to reduce the memory costs. We adapt the Min-Block Greedy approach [29] and the Round-Robin Greedy approach [3] to the constrained case. To the best of our knowledge, this work is the first (as far as we know) to study the robust submodular partitioning problem under all of the various constraints (cardinality, intersection of matroid, knapsack). Wang et al. [25] is a special case of our work as it only studies the cardinality constraint. Cotter et al. [8] studies (as well as allowing multiple blocks to be jointly scored) a matroid constrained "groupings" (i.e.,coverings, packings, or partitions) problem but only a fractional subset of groups (rather than the minimum of the groups), is guaranteed to have values larger than the bounded max-min OPT, while our bound compares the min block evaluation to the optimal max-min value. Another line of related research is the submodular load balancing problem, which minimizes the maximum-valued block in the partition according to the submodular evaluations. In contrast to promoting diversity of each block for the robust submodular partition problem, submodular load balancing encourages every block to contain redundant items, similar to standard clustering objectives. Theoretically, this problem has been shown to be much harder as Svitkina & Fleischer [22] shows a information theoretical lower bound of $o(\sqrt{\frac{n}{\log n}})$, and also gives a sampling algorithm to match the lower bound up to constant factors. Similar to the max-min case, Ghodsi et al. [11] uses the ellipsoidal approximation to get a bound of $\mathcal{O}(\sqrt{n}\log n)$. Wei et al. [29] gives a Lovász extension based relaxation algorithm and achieves a bound of m.

3 Preliminaries and Formulation

With a ground set V of n items, a submodular function f is a set function $2^V \to \mathbb{R}$ that satisfies the property: $f(A)+f(B) \geq f(A\cup B)+f(A\cap B)$, where $A, B \subseteq V$. Equivalently, a submodular function is characterized by diminishing returns: $f(v|A) \geq f(v|B)\ \forall v \notin B$ and $A \subset B \subseteq V$, where $f(v|B) = f(\{v\}\cup B) - f(B)$. Submodular functions naturally describe the diversity or representativeness of a given set of items. Many simple greedy-based algorithms have been developed to solve optimization problems involving submodular functions, giving both theoretical approximation guarantees, as well as good empirical performance. We restrict the submodular functions discussed in this paper to be monotone non-decreasing and normalized, i.e., $f(B) \geq f(A)\ \forall A \subseteq B \subseteq V$, $f(\emptyset) = 0$.

A matroid $\mathcal{M} = (V, \mathcal{I})$ is a set system that describes the independence relationships among the subsets of the ground set V. $\mathcal{I}$ is a set of subsets of V and every $S \in \mathcal{I}$ is considered an independent subset. The matroid rank function is defined as $r_{\mathcal{M}}(A) = \max\{|S| : S \subseteq A, S \in \mathcal{I}\}$. $r_{\mathcal{M}}(V)$ indicates the maximum size of a subset that may be independent according to the matroid $\mathcal{M}$. All subsets of cardinality $\leq k$ with some integer $k > 0$ form a *uniform matroid*, which we denote by $\mathcal{M}_k^u$. A partition matroid is one where we partition the ground set into blocks, and a set is independent if it intersects each block by no more than a block-specific limit. We define a particularly useful partition matroid on an expanded ground set $\bar{V}$ as follows: We first duplicate the ground set m times, creating $V_1 = V_2 = \ldots V_m = V$, which are ground set copies. We create an expanded ground set $\bar{V} = \uplus_{j=1:m} V_j$ as the disjoint union. A subset $S \subseteq \bar{V}$ is independent in $\mathcal{M}_m^p$ if for every element $v \in V$, let its m copies in $\bar{V}$ be $\{v_1, v_2, \ldots, v_m\}$, we have $|S \cap \{v_1, v_2, \ldots, v_m\}| \leq 1$, i.e., S contains at most one copy of element v. Apart from the uniform matroid and this particular partition matroid, there are many other matroids reflecting a natural notion of independence, for example, the linearly-independent set of real vectors and the spanning trees in a graph. In the below, we use both $S \in \mathcal{M}$ and, when clear, $S \in \mathcal{I}$, to indicate that S is independent in the matroid $\mathcal{M} = (V, \mathcal{I})$.

Matroids are often used as constraints in submodular optimization problems: $\max_{S\in\mathcal{I}} f(S)$ with a matroid $\mathcal{M} = (V, \mathcal{I})$. When $\mathcal{M}$ is a uniform matroid $\mathcal{M}_k^u$, this reduces to the cardinality submodular max and the greedy algorithm gives a $1-e^{-1}$ bound [9]. For a general constraint with the intersection of p matroids, the simple greedy algorithm gives a $\frac{1}{p+1}$ bound [10]. Suppose we represent a set S as a binary indicator vector $x_S \in \{0,1\}^n$, i.e., $\forall i \in [n], x_S[i] = 1$ if $v_i \in S$ or otherwise $x_S[i] = 0$. Then for all the independent sets of a matroid $\mathcal{M} = (V, \mathcal{I})$, the convex hull over all the $x_S, S \in \mathcal{I}$ forms a polytope, which is called the matroid polytope $\mathcal{P}_{\mathcal{M}}$ of matroid $\mathcal{M}$ [9]. Based on the convex property of the matroid polytope, algorithms [4–6, 24] have been proposed to firstly solve a continuous extension of the submodular optimization problem under the matroid polytope constraint, which generates a fractional solution in $[0,1]^n$, and then round the fractional solution to an integral solution to get the resulting set. The continuous greedy algorithm [4] gives a $1-e^{-1}$ guarantee under a single

matroid constraint using pipage rounding [1, 4]. Interestingly, running the continuous greedy under a partition matroid constraint (submodular welfare problem) gives a uniform fractional solution, i.e., on the expanded ground set $\bar{V}$, the fractional solution $x = (\frac{1}{m}, \frac{1}{m}, \ldots, \frac{1}{m})$ (i.e., assigning $\frac{1}{m}$ of every element to each block) leads to a $1 - e^{-1}$ bound in expectation for an integral solution that assigns each element in V uniformly to one of the m blocks. This observation also constitutes the basic idea of Round-Robin Greedy for solving the robust submodular partition problem [3], which we will discuss in more detail later.

For a submodular function f on a ground set V, the robust submodular partition problem (submodular fair allocation) [12] is defined as:

$$\max_{\pi \in \Pi(V,m)} \min_{A \in \pi} f(A), \tag{1}$$

where m is the number of blocks in a partition, we denote all possible partitions with m blocks of ground set V as $\Pi(V, m)$, and one partition π with $|\pi| = m$ is a collection of m disjoint sets. Equivalently, we can represent the partition using a partition matroid constraint on the expanded ground set $\bar{V}$:

$$\max_{S \subseteq \bar{V}, S \in \mathcal{M}_m^p} \min_{j \in [m]} f(S \cap V_j). \tag{2}$$

Intuitively, the above optimization for robust submodular partitioning encourages the minimum-valued block to have a high submodular evaluation. Compared to the submodular welfare problem, the robust submodular partition promotes fairness for every one of the partition blocks.

There have been three recent approximation algorithms developed to solve Eq. (1). Particularly, Wei et al. [29] uses a Min-Block Greedy algorithm, which greedily adds the element with the largest gain to the block with the minimum evaluation. Barman & Krishna Murthy [3] propose a Round-Robin Greedy algorithm, which iteratively traverses all the blocks in a fixed order, and greedily adds an element with the largest gain to each block. Ghodsi et al. [11] applies a local search approach, which starts with an arbitrary partition and keeps moving an element from a non-minimum block to the minimum block if this relocation improves the objective by certain threshold until no such element can be found.

For [3] and [11], they both require guessing the optimal solution's value, and they need to run multiple instances of their algorithms with the guessed optimal values as an exponentially decreasing sequence from the maximal possible value $f(V)$ to the optimal solution value $\mu = \max_{\pi \in \Pi(V,m)} \min_{S \in \pi} f(S)$. With the exponential decreasing factor as $1 + \delta$, the running time (in terms of submodular function calls) is $O(n^2 \frac{1}{\delta} \log \frac{f(V)}{\mu})$ for [3], and $O(n^2 m^2 \frac{1}{\delta} \log \frac{f(V)}{\mu})$ for [11]. Min-Block Greedy, a much simpler algorithm, has a running time of $O(n^2)$. Note that the settings of [3] and [11] are slightly more general than Eq. (1) as the submodular function for each block can be different. But it's not the heterogeneous case either as they focus on a different notion of optimality.

3.1 Discussions about the Optimality in [3, 11] and the Heterogeneous Case

The works [3, 11] both study the partitioning problem in the economics context of fair allocation of indivisible goods. In such a setting, every block is an agent, and we want to find an allocation of the goods to each agent in a fair manner, so that each agent's evaluation of the allocated goods to himself is optimized. Each agent can have different evaluations for the goods, which means that the submodular function for each block can be different. Taking [3] as an example, the theoretical guarantee they prove is

Lemma 1 (Theoretical Guarantee in [3])**.** *Let $A_1, A_2, \ldots, A_m$ be the solution to the Round-Robin Greedy algorithm for the unconstrained problem, and given m submodular functions for the m blocks as f_j for $j = 1, 2, \ldots m$, for every agent j we have $f_j(A_j) \geq \frac{1-e^{-1}}{3} \max_{\pi \in \Pi(V,m)} \min_{S \in \pi} f_j(S)$.*

Intuitively, the bound guarantees that based on each agent's evaluation f_j, the goods allocated to himself is not bad compared to the worst block in the allocation. When all the f_j's are the same, the bound reduces to the bound for the homogeneous case of robust submodular partitioning, which is the focus of this paper. The heterogeneous case for robust submodular partitioning is different, as it requires to show a bound like (suppose the algorithm solution is A_j for $j \in [m]$)

$$\min_{j \in [m]} f_j(A_j) \geq \gamma \max_{\pi \in \Pi(V,m)} \min_{S \in \pi} f_j(S). \tag{3}$$

We give an example of the m functions so that the two guarantees vary. Say we have a predefined partition over the ground set V into m blocks as $C_1, C_2, \ldots, C_m$ and $\cup_{j\in[m]} C_j = V$, $C_j \cap C_{j'} = \emptyset$, $|C_j| = |C_{j'}|$ (assume $|V|$ is a multiple of m). Let $f_j(S) = |S \cap C_j|$. The optimal solution to the heterogeneous case $\max_{\pi\in\Pi(V,m)} \min_{j\in[m]} f_j(A_j)$ is to assign the items in the same way as the predefined partition $C_1, C_2, \ldots C_m$, and the optimal solution value is $|C_j|$. However, for the bound in [3, 11], the optimal solution $\{O_1, O_2, \ldots, O_m\} \in \operatorname{argmax}_{\pi\in\Pi(V,m)} \min_{A\in\pi} f_j(A)$ is to intersect each predefined block equally, i.e., $|O_j \cap C_{j'}| = \frac{|C_{j'}|}{m}$ $\forall j, j' \in 1, \ldots, m$. Therefore, the optimal solution value is then $\frac{|C_j|}{m}$.

In this work, we study the constrained case for the submodular robust partition problem in the homogeneous setting (submodular functions for all blocks are the same). We extend Min-Block Greedy algorithm [29] and Round-Robin Greedy algorithm [3] to adapt to various constraints, e.g., cardinality, matroid, and intersection of matroids.

4 Min-Block Greedy Based Algorithms

Wei et al. [29] proposes a Min-Block Greedy Algorithm 2 for Eq. (1), which loops over n iterations, and at every iteration, for the minimum-valued block $A_{j^*} \in \operatorname{argmin}_j f(A_j)$, it finds the element with the largest gain $f(v|A_{j^*})$. Wei et al. [29] proves a $1/m$ bound of Min-Block Greedy. In fact, their proof works for a simpler algorithm, Min-Block Streaming Algorithm 1, which assumes that the algorithm accesses elements from the ground set in an arbitrary order as a stream $V = (v_1, v_2, \ldots, v_n)$, and it assigns the incoming element to the block with the least evaluation. We denote the optimal partition to Eq. (1) as $\pi^* = \{O_1, O_2, \ldots, O_m\}$.

Lemma 2 (**Unconstrained Min-Block Streaming[29]**). *For a ground set V and its elements $(v_1, v_2, \ldots, v_n)$ coming in an arbitrary streaming order, the output solution of Alg. 1 has $\min_{j\in[m]} f(A_j) \geq \frac{1}{m} \min_{j\in[m]} f(O_j)$.*

Corollary 1 (**Unconstrained Min-Block Greedy[29]**). *The output solution of Alg. 2 has $\min_{j\in[m]} f(A_j) \geq \frac{1}{m} \min_{j\in[m]} f(O_j)$ since the order of adding elements in Min-Block Greedy is one possible order of the ground set elements.*

Intuitively, Alg. 2 optimizes the objective Eq. (1) greedily, i.e., it always increases the current value (the minimum-block evaluation) with the largest possible gain, while the performance of Alg. 1 greatly depends on the order of elements, so it might seem that the bound for Min-Block Greedy should improve upon the current $\frac{1}{m}$ bound. However, as shown in our new result in the following lemma, the bound in Corollary 1 is tight.

Lemma 3 (**Tightness of Corollary 1).** *$\forall \epsilon > 0$, $\exists$ a submodular function f such that the output solution of Alg 2 $\min_{j=1:m} f(A_j) = \frac{1}{m} \min_{j=1:m} f(O_j) + \epsilon$.*

Algorithm 1: Min-Block Streaming

input : submodular function f, ground set as a stream $V = (v_1, v_2, \ldots, v_n)$, number of blocks m

$R := V$;
Let $A_1 = A_2 = \ldots = A_m = \emptyset$;
for $i = 1 : n$ **do**
  $j^* \in \operatorname{argmin}_j f(A_j)$;
  $A_{j^*} := A_{j^*} \cup \{v_i\}$;
return $(A_1, A_2, \ldots, A_m)$

Algorithm 2: Min-Block Greedy

input : submodular function f, ground set V, number of blocks m

$R := V$;
Let $A_1 = A_2 = \ldots = A_m = \emptyset$;
while $R \neq \emptyset$ **do**
  $j^* \in \operatorname{argmin}_j f(A_j)$;
  $v^* \in \operatorname{argmax}_{v\in R} f(v|A_{j^*})$;
  $A_{j^*} := A_{j^*} \cup \{v^*\}$;
  $R := R \setminus \{v^*\}$;
return $(A_1, A_2, \ldots, A_m)$

We elaborate on how to construct the submodular function in Appendix A. The key idea is that we can find a set-cover function where even though Min-Block Greedy selects the element with the largest gain, the element can still be quite redundant with the current minimum block. Say the current minimum block is A, and the maximum-gain element is v chosen by the greedy step, meaning $f(v|A)$ is larger than $f(v'|A)$ for $v' \in R \setminus v$. However, $\frac{f(v|A)}{f(v)}$ can still be very small, i.e., the area covered by v according to the set-cover function is already mostly covered by A. On the

other hand, the optimal solution can fully utilize $f(v)$ thus making a more lonesome v cover a much larger area overall. Note that Lemma 3 also serves as the tightness for Lemma. 2 since the order of adding elements in Min-Block Greedy follows a streaming order.

More generally, given a constraint $\mathcal{C}$, we define the constrained robust submodular partition as:

$$\max_{\pi \in \Pi(V,m,\mathcal{C})} \min_{A \in \pi} f(A), \tag{4}$$

where $\Pi(V, m, \mathcal{C})$ is the set of all possible partitions on set V into m blocks such that for every partition $\pi \in \Pi(V, m, \mathcal{C})$, every block $A \in \pi$ should satisfy the constraint $A \in \mathcal{C}$. We denote the optimal partition in Eq. (4) as $\pi^*_C = \{O^{\mathcal{C}}_1, O^{\mathcal{C}}_2, \ldots, O^{\mathcal{C}}_m\}$. We remark that due to the constraints, some elements might not be assigned to a partition block, so strictly speaking the solution is an allocation (or "grouping", Cotter et al. [8]) of elements rather than a partition.

For now, we will take $\mathcal{C}$ as any down-closed constraint: Let $\mathcal{C}$ be a collection of subsets of the ground set V, and by satisfying the constraint, we require the solution A to be one of the subsets in $\mathcal{C}$. The down-closed property means that if $A \in \mathcal{C}$ we have $B \in \mathcal{C}$ for any $B \subseteq A$. Following Eq. (4), we can define the constrained problem in terms of the expanded subset $\bar{V}$:

$$\max_{S \subseteq \bar{V}, S \in \mathcal{M}^p_m, \forall j: (S \cap V_j) \in \mathcal{C}} \min_{j \in [m]} f(S \cap V_j). \tag{5}$$

Based on the Min-Block Greedy algorithm for the unconstrained case, we propose a natural extension to the constrained case (Alg. 3), where at every iteration, for the minimum-valued block A_{j^*}, we greedily find the best element v^* that retains block feasibility under the constraint $\mathcal{C}$, i.e., $\{v^*\} \cup A_{j^*} \in \mathcal{C}$. If we cannot find any element in the remaining set to add to the current minimum block, we remove the current block from the candidate blocks and move to the next smallest valued block.

In Line 6 of Alg. 3, we call a subroutine GreedyStep$(R, \mathcal{C}, A_{j^*})$ to greedily find a feasible element. The subroutine varies according to the type of constraint $\mathcal{C}$. Particularly, for the constrained submodular maximization problem defined as

$$\max_{S \subseteq V, S \in \mathcal{C}} f(S). \tag{6}$$

GreedyStep$(\cdot)$ is shared by Alg. 3 and Alg. 4, and if Alg. 4 is an approximation algorithm of solving Eq. (6) with some bound α, we can prove the following result for Alg. 3.

Algorithm 3: Constrained Min-Block Greedy

input : submodular function f, ground set V, number of blocks m, constraint $\mathcal{C}$

1 Let $A_1 = A_2 = ... = A_m = \emptyset$;
2 Let $J = [m]$, $R = V$;
3 **while** $R \neq \emptyset$ *and* $J \neq \emptyset$ **do**
4 $\quad j^* \in \text{argmin}_{j \in J} f(A_j)$;
5 $\quad$ **if** $\exists v \in R$ *s.t.* $A_{j^*} \cup \{v\} \in \mathcal{C}$ **then**
6 $\quad\quad v^*$:= GreedyStep$(R, \mathcal{C}, A_{j^*})$;
7 $\quad\quad A_{j^*} := A_{j^*} \cup \{v^*\}$, $R := R \setminus \{v^*\}$;
8 $\quad$ **else**
9 $\quad\quad a^* \in \text{argmax}_{a \in A_{j^*} \cup R} f(\{a\})$;
10 $\quad\quad$ **if** $f(\{a^*\}) \geq f(A_{j^*})$ **then**
11 $\quad\quad\quad A_{j^*} := \{a^*\}$, $R := R \setminus \{a^*\}$;
12 $\quad\quad$ Let $J = J \setminus j^*$;
13 **return** $(A_1, A_2, ..., A_m)$

Algorithm 4: Constrained Submodular Greedy Max

input : submodular function f, ground set V, constraint $\mathcal{C}$

$R := V$;
Let $S^g = \emptyset$;
while $R \neq \emptyset$ **do**
$\quad$ **if** $\exists v \in R$ *s.t.* $A_{j^*} \cup \{v\} \in \mathcal{C}$ **then**
$\quad\quad v^*$:= GreedyStep$(R, \mathcal{C}, S^g)$;
$\quad\quad S^g := S^g \cup \{v^*\}$;
$\quad\quad R := R \setminus \{v^*\}$;
$\quad$ **else**
$\quad\quad$ Break;
$a^* \in \text{argmax}_{a \in V} f(\{a\})$;
return $\text{argmax}_{A \in \{S^g, \{a^*\}\}} f(A)$

Theorem 1 (Constrained Min-block Greedy). *Given a constraint $\mathcal{C}$, if the greedy solution S^g to problem $\max_{S \in \mathcal{C}} f(S)$ using Alg. 4 has a bound of α, i.e., $f(S^g) \geq \alpha \max_{S \in \mathcal{C}} f(S)$, then the solution of Alg. 3 has $\min_{j \in [m]} f(A_j) \geq \frac{\alpha}{\alpha m + 1} \min_{j \in [m]} f(O^{\mathcal{C}}_j)$. Assuming GreedyStep$(\cdot)$ takes $\mathcal{O}(1)$ oracle calls, the time complexity of of Alg. 4 is $\mathcal{O}(n^2)$.*

The general idea of the proof (details in Appendix A) is that we divide the ground set V into two disjoint parts $V = V' \cup R'$, where the min block in the output solution A intersecting V' corresponds

to the min block solution of an instance of unconstrained robust partition problem (Eq. (1)) defined on the ground set V', and $A \cap R'$ with the other part corresponds to the solution of an instance of submodular maximization (Eq. (6)) under the constraint $\mathcal{C}$ defined on the ground set $A \cup R'$. We bound the two parts separately and combine them to obtain the above bound.

Corollary 2 (**Cardinality Constrained Min-block Greedy**). *For $\mathcal{C}$ as a cardinality constraint, the output of Alg 3 has* $\min_{j=1:m} f(A_j) \geq \frac{1}{m+\frac{1}{1-e^{-1}}} \min_{j=1:m} f(O_j^{\mathcal{C}})$.

Corollary 3 (**Matroid Constrained Min-block Greedy**). *For $\mathcal{C}$ as an intersection of p matroids constraint, the output of Alg 3 has* $\min_{j=1:m} f(A_j) \geq \frac{1}{m+p+1} \min_{j=1:m} f(O_j^{\mathcal{C}})$.

Corollary 4 (**Knapsack Constrained Min-block Greedy**). *For $\mathcal{C}$ as a knapsack constraint, the output of Alg 3 has* $\min_{j=1:m} f(A_j) \geq \frac{1}{m+\frac{2}{1-1/e}} \min_{j=1:m} f(O_j^{\mathcal{C}})$.

For $\mathcal{C}$ as a cardinality constraint, GreedyStep$(\cdot)$ just picks the element with the largest gain assuming the block has not yet reached the cardinality limit k. For $\mathcal{C}$ as an intersection of p matroid constraints, GreedyStep$(\cdot)$ finds the element v^* that has the largest gain $f(v^*|A_{j^*})$ assuming the block can be kept feasible, i.e., $v^* \cup A_{j^*} \in \mathcal{C}$. For $\mathcal{C}$ as a knapsack constraint with the weight of each element v as $w(v)$, GreedyStep$(\cdot)$ finds the element v^* with the largest ratio $\frac{f(v^*|A_{j^*})}{w(v)}$ assuming the sum of weights can be kept below the given budget. In line 9-11 of Alg. 3 and line 10-11 of Alg. 4, we include an extra step of comparing with the largest singleton value. Such step is redundant when $\mathcal{C}$ is an intersection of matroid constraints, but is essential for the knapsack constraint case, as the modified greedy algorithm for the knapsack problem [16] requires this extra step or otherwise α is unbounded. Due to the tightness of the $\frac{1}{m}$ bound we have proved for the unconstrained case, the $\frac{1}{m}$ dependence in the constrained bound cannot be improved.

5 Round-Robin Greedy Based Algorithms

Barman & Krishna Murthy [3] propose a round-robin style algorithm for the unconstrained robust submodular partition problem (Eq. (1)) and gives a constant bound of $\frac{1-e^{-1}}{3}$ with weakly polynomial running time. Compared to Min-Block Greedy, Round-Robin Greedy requires guessing the optimal values by an exponentially decreasing sequence, and for each guessed value, it runs one instance of the round-robin subroutine. Specifically, suppose $\mu = \min_{j\in[m]} f(O_j)$, i.e., μ is the optimal solution value for the unconstrained case, then for a parameter $\delta > 0$, Round-Robin Greedy runs the round-robin subroutine with the guessed optimal values from a sequence $(f(V), \frac{f(V)}{1+\delta}, \frac{f(V)}{(1+\delta)^2}, \ldots)$ and ends when the guessed value is no larger than μ. The running time of each round-robin subroutine is $\mathcal{O}(n^2)$, as it greedily finds the element with the largest gain by iterating over all the remaining elements. There are $\log_{1+\delta} \frac{f(V)}{\mu}$ guessed values in the exponentially decreasing sequence, so the overall running time is $\mathcal{O}(n^2 \log_{1+\delta} \frac{f(V)}{\mu}) = \mathcal{O}(n^2 \frac{1}{\delta} \log \frac{f(V)}{\mu})$. Note that since we use a $(1+\delta)$ factor exponentially decreasing sequence, we lose a $(1+\delta)$ factor in the approximation bound, which can be improved arbitrarily by using a smaller δ value but with a cost of running more instances of the round-robin subroutine.

The major idea behind Round-Robin Greedy comes from the solution of Continuous Greedy for the submodular welfare problem, which is an uniform fractional vector $x = (\frac{1}{m}, \frac{1}{m}, \ldots, \frac{1}{m})$ with the length of x equal to the size of the expanded ground set $|\bar{V}|$. Let F be the multilinear extension [4] of f, Continuous Greedy gives a bound that $F(x) = \mathbb{E}_{R\sim x} f(R) \geq (1-e^{-1}) \max_{\pi\in\Pi(V,m)} \sum_{A\in\pi} f(A)$, where $\mathbb{E}_{R\sim x} f(R)$ takes the expectation of $f(R)$ on a random set R with each element sampled independently according to the probability in the fractional vector x. Note that the hardness for submodular optimization under a matroid constraint is $1-e^{-1}$, which means that the random assignment strategy achieves the best possible theoretical bound on the submodular welfare problem.

Round-Robin Greedy can be thought as a rounding mechanism for the fractional solution x. The round-robin style iteration is similar to the uniform random assignment in a deterministic manner, and by greedily finding the element, the value of every block can be bounded against that for the random assignment. In fact, Round-Robin Greedy bounds every block A_j to be $f(A_j) \geq \frac{1}{3}\frac{F(x)}{m}$, and since the welfare solution bounds the robust solution in terms of the sum: $\max_{\pi\in\Pi(V,m)} \sum_{A\in\pi} f(A) \geq \sum_{j\in[m]} O_j \geq m\mu$, we get the desired bound for the robust partition problem.

We extend Round-Robin Greedy to the constrained case (Eq. (4)) firstly with $\mathcal{C}$ as a cardinality constraint k. This is a relatively simple case due to the nature of Round-Robin Greedy that every block gets assigned with the same number of elements at the end of every round-robin iteration. We present the modified algorithm in Alg. 5, which also helps to explain the essential ideas of the original Round-Robin Greedy as we describe below.

Lemma 4 (**Cardinality Constrained Round-Robin).** *For the problem in Eq.* (4)*, with $\mathcal{C}$ as a cardinality constraint k, Alg. 5 gives a solution* $\min_{j\in[m]} f(A_j) \geq \frac{(1-e^{-1})^2}{3} \min_{j\in[m]} f(O_j^k)$.

Algorithm 5: Cardinality Round-Robin Greedy

input : f, V, m, cardinality constraint k, discounting factor for guessing optimal δ

1 Let τ be the solution value of Alg. 3;
2 Let $high = \lceil \log_{1+\delta}(m+2) \rceil$, $low = 0$;
3 Create a sequence of guessed values: $(\tau, (1+\delta)\tau, (1+\delta)^2\tau, \ldots, (1+\delta)^{high}\tau)$;
4 Create an empty solution ($\emptyset$ for each block in the partition) for each guessed value $\pi_0, \pi_1, \ldots, \pi_{high}$;
5 **while** $high \geq low$ **do**
6 Let $idx = \lfloor (high + low)/2 \rfloor$; Let $A_1 = A_2 = \ldots = A_m = \emptyset$;
7 Let $V' = \{v | v \in V, f(v) \leq \frac{(1-e^{-1})^2}{3}(1+\delta)^{idx}\tau\}$; Let $G = V \setminus V'$;
8 Assign G to $A_{m-|G|+1}, A_{m-|G|+2}, \ldots, A_m$ with one element per block;
9 Let $m' = m - |G|$;
10 Let $A'_1, A'_2, \ldots, A'_{m'}$ be the solution to $\max_{\pi\in\Pi(V',m',k)} \sum_{S\in\pi} f(S)$ using continuous greedy and swap rounding; Let $V'' = \cup_{j\in[m]'} A'_j$;
11 Let $\{A_1, A_2, \ldots, A_{m'}\} = RR(f, V'', m', \mathcal{M}_k^u, [m'])$;
12 **if** $f(A_j) \geq \frac{(1-e^{-1})^2}{3}(1+\delta)^{idx}\tau \; \forall j \in [m']$ **then**
13 Let $\pi_{idx} = \{A_1, A_2, \ldots, A_m\}$; Let $low = idx + 1$;
14 **else**
15 Let $high = idx - 1$;
16 **return** best of $\pi_0, \pi_1, \ldots, \pi_{high}$;

Here is how we achieve strongly-polynomial time. Different from the original Round-Robin Greedy, which performs a grid search over guessed optimal values, we perform a binary search over the sequence of values and therefore the number of outer iterations is reduced. Most importantly, we use the Min-Block Greedy solution's value as the minimum guessed value τ. Because of the $\frac{1}{m+\frac{1}{1-1/e}}$ bound of the Min-Block Greedy solution, the maximum guessed value is thus bounded by $(m+2)\tau$. We then create a $1+\delta$-factor exponential decreasing sequence between τ and $(m+2)\tau$ to binary search for the optimal solution value. This improves the number of outer iterations of the algorithm to $\mathcal{O}(\log\log_{1+\delta} m) = \mathcal{O}(\log\log m + \log\frac{1}{\delta})$, which is strongly-polynomial while the number of outer iterations $\mathcal{O}(\log_{1+\delta}\frac{f(V)}{\mu})$ for the original unconstrained case is only weakly-polynomial as it has a $\log$ dependence on the function value.

Algorithm 6: Round-Robin Greedy Iterations ($RR(f, R, m', \mathcal{M}, J)$)

input : f, R, m', matroid constraint $\mathcal{M}$, set of block indices J

while $J \neq \emptyset$ *and* $R \neq \emptyset$ **do**
  for $j \in [m']$ **do**
    if $j \in J$ **then**
      if $\exists v \in R$ *s.t.* $A_j \cup \{v\} \in \mathcal{M}$ **then**
        $v^* \in \operatorname{argmax}_{v\in R, A_j\cup v\in\mathcal{M}} f(v|A_j)$;
        $A_j := A_j \cup \{v^*\}$;
        $R := R \setminus \{v^*\}$;
      else
        Let $J = J \setminus j$;
return $(A_1, A_2, \ldots, A_{m'})$

In every outer iteration (Line 12-15), Alg 5 checks if the round-robin solution based on the guessed optimal value $(1+\delta)^{idx}\tau$ satisfy the approximation bound, i.e., $f(A_j) \geq \frac{(1-1/e)^2}{3}(1+\delta)^{idx}\tau$ $\forall j \in [m]'$. If the bound is (not) satisfied, the guessed value is large (small) and we move to an increased (decreased) search value. Within every outer iteration, we perform round-robin greedy (iterate over every block in some fixed order and greedily add to the block the element with the largest gain). Line 10 of Alg. 5 is the major change to Round-Robin Greedy specifically for the cardinality constraint case, where we first find the solution to the cardinality constrained submodular

welfare problem $\max_{\pi \in \Pi(V', m', k)} \sum_{S \in \pi} f(S)$, and then only apply Round-Robin Greedy to the union V'' of the solution $A'_1, A'_2, \ldots, A'_{m'}$.

The running time of Alg. 5 is similar to Round-Robin Greedy, with additional costs caused by Line 10, which solves a cardinality constrained submodular welfare problem. Using Continuous Greedy and swap rounding [7] for Line 10 can be quite costly ($\mathcal{O}(n^5)$ for the inner loop), which may improve in the future by a better algorithm. In Alg. 7, we propose another algorithm that addresses the constrained robust submodular problem with $\mathcal{C}$ as any matroid constraint $\mathcal{M}$ and incurs no additional computation costs compared to Round-Robin Greedy.

Theorem 2 (**Matroid Constrained Round-Robin).** *For the problem in Eq.* (4)*, with $\mathcal{C}$ as any matroid constraint $\mathcal{M}$, Alg 7 gives a solution* $\min_{j \in [m]} f(A_j) \geq \frac{(1-e^{-1})}{5} \min_{j \in [m]} f(O_j^{\mathcal{M}})$. *The time complexity of of Alg. 7 is* $\mathcal{O}(n^2(\log\log m + \log \frac{1}{\delta}))$.

Algorithm 7: Matroid Round-Robin Greedy

input : f, V, m, matroid constraint $\mathcal{M}$, discounting factor for guessing optimal δ

1 Let τ be the solution value of Alg. 3;
2 Let $high = \lceil \log_{1+\delta}(m+2) \rceil$, $low = 0$;
3 Create a sequence of guessed values: $(\tau, (1+\delta)\tau, (1+\delta)^2\tau, \ldots, (1+\delta)^{high}\tau)$;
4 Create an empty solution ($\emptyset$ for each block in the partition) for each guessed value $\pi_0, \pi_1, \ldots, \pi_{high}$;
5 **while** $high \geq low$ **do**
6 Let $idx = \lfloor (high + low)/2 \rfloor$; Let $A_1 = A_2 = \ldots = A_m = \emptyset$;
7 Let $V' = \{v | v \in V, f(v) \leq \frac{1-e^{-1}}{5}(1+\delta)^{idx}\tau\}$; Let $G = V \setminus V'$;
8 Assign G to $A_{m-|G|+1}, A_{m-|G|+2}, \ldots, A_m$ with one element per block;
9 Let $m' = m - |G|$;
10 Let $\{A_1, A_2, \ldots, A_{m'}\} = RR(f, V', m', \mathcal{M}, [m'])$;
11 **if** $f(A_j) \geq \frac{(1-e^{-1})}{5}(1+\delta)^{idx}\tau \; \forall j \in [m']$ **then**
12 Let $\pi_{idx} = \{A_1, A_2, \ldots, A_m\}$; Let $low = idx + 1$;
13 **else**
14 Let $high = idx - 1$;
15 **return** best of $\pi_0, \pi_1, \ldots, \pi_{high}$;

Comparing to Alg. 5, the major change in Alg. 7 is (1) we do not need to run the costly Continuous Greedy and swap rounding to get a solution to the constrained welfare problem, which makes the algorithm applicable in practice; (2) for the $RR(\cdot)$ subroutine, we find a feasible element and add it to the current block, and we remove a block from the candidate set J if there are no element in the remaining set that can be added to the block without violating the matroid constraint. Note for the cardinality constraint case, we can always find a feasible element until every block has k elements. The overall running time is $\mathcal{O}(n^2(\log\log m + \log \frac{1}{\delta}))$. The general idea of proving Theorem 2 is to bound the solution to the fractional solution of the continuous relaxation. For every block in the solution, we inspect the elements that have been evaluated during the greedy step. For those elements with large gains when being evaluated but not added due to the violation of the matroid constraint, we bound their gains as submodular maximization with a matroid constraint on a reduced ground set. For the remaining elements, we bound their gains by the greedy step and together we get the desired bound.

6 Experiments

We empirically test Algs. 3 and. 7 on the CIFAR-10 training set [19] ($|V|$ = 50000). We use facility location as our submodular function, i.e., $f(S) = \sum_{v \in V} \max_{v' \in S} \text{sim}(v, v')$, where $\text{sim}(v, v')$ measures the affinity between elements v and v'. This function is widely used and naturally describes a subset's diversity via its similarities/distances to all other points in the ground set, and it has achieved much practical success [27, 26, 23]. For similarity, we use a Gaussian kernel with L2 distances, i.e., $\text{sim}(v, v') = exp(\frac{-||v-v'||_2}{\sigma})$, where σ is the bandwidth of the kernel, set to the average L2 distance, i.e., $\sigma = \sum_{v,v' \in V} ||v - v'||_2 / n^2$. The features used to calculate the L2 distance is the bottleneck layer's outputs generated by a deep auto-encoder model (details are in Appendix D).

We test the algorithms and compare their objective values (Eq. (4)) with a matroid constraint where we limit the number of samples selected for each class in CIFAR-10 for each block (CIFAR-10 has 10 classes). We compare the two algorithms with a random selection baseline — we randomly sample from each class and assign to each block with the constrained number of elements. Hence, the random selection results satisfy the constraints. In Fig. 1, we report the results for different matroid constraints with various block sizes. The random baseline results are reported with means and standard-deviations over 10 runs. For all cases, we see that both Alg. 3 and Alg. 7 significantly outperform the baselines. Although Alg. 7 has a better theoretical bound, Alg. 3 consistently gives better performance. Intuitively, Alg. 3 directly optimizes the objective as it greedily adds elements to the minimum-valued block. We expect Alg. 3 to perform better in practice compared to Alg. 7 as Alg. 7 has a fixed ordering of the blocks, and the minimum-valued block tends to be the last block in the ordering, in which case it does not get to select samples that have already been selected by prior blocks in the ordering. In Fig. 2, we use the partitioned blocks as minibatches to train a ResNet-9 model and compare their performance on the test set. We observe that the blocks with higher submodular evaluations tend to generate better performance for the trained model. We also provide results on synthetic data in Appendix Section C.

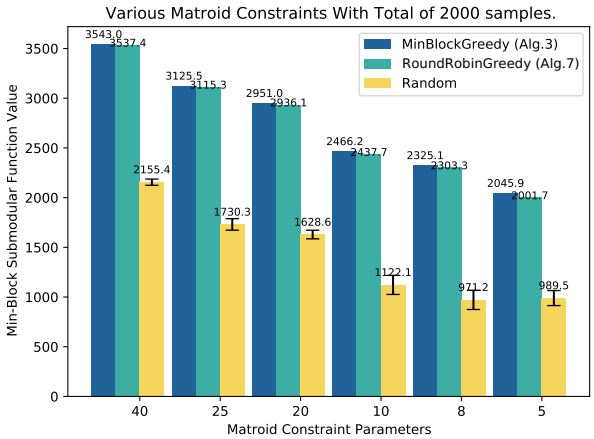

Figure 1: Matroid Constrained results. The x-axis denotes the number of samples per class for the matroid. A total of 2000 samples are selected for each case.

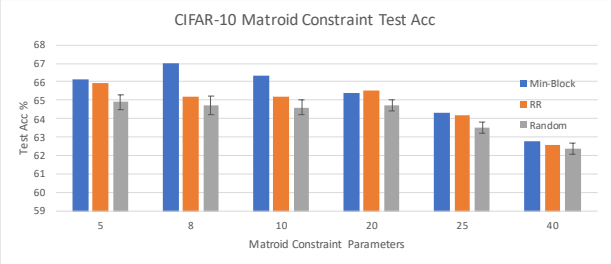

Figure 2: Training ResNet-9 (Myrtle AI) on partitioned mini-batches under matroid constraints with various parameters. The y-axis denotes the test set accuracy. The x-axis denotes the number of samples per class for the matroid.

7 Conclusions

We study the problem of constrained submodular robust partitioning. We propose two classes of algorithms, Min-Block Greedy and Round-Robin Greedy based, and prove approximation bounds under various constraints. This improves the applicability of the robust partitioning framework to different scenarios. In future work, we wish to extend the current approach to the heterogeneous submodular partitioning setting where each block may be evaluated by a different submodular function. Given the good performance of Alg. 3 in practice, it is worth investigating if further conditions or modifications to Alg. 3 yield improved theoretical bounds.

This work was supported in part by the CONIX Research Center, one of six centers in JUMP, a Semiconductor Research Corporation (SRC) program sponsored by DARPA.

References

[1] Alexander A Ageev and Maxim I Sviridenko. Pipage rounding: A new method of constructing algorithms with proven performance guarantee. *Journal of Combinatorial Optimization*, 8(3): 307–328, 2004.

[2] Arash Asadpour and Amin Saberi. An approximation algorithm for max-min fair allocation of indivisible goods. In *SICOMP*, 2010.

[3] Siddharth Barman and Sanath Kumar Krishna Murthy. Approximation algorithms for maximin fair division. In *Proceedings of the 2017 ACM Conference on Economics and Computation*, pp. 647–664, 2017.

[4] G. Calinescu, C. Chekuri, M. Pal, and J. Vondrák. Maximizing a monotone submodular function under a matroid constraint. *IPCO*, 2007.

[5] C. Chekuri, J. Vondrák, and R. Zenklusen. Submodular function maximization via the multilinear relaxation and contention resolution schemes. *Proceedings of the 43rd annual ACM symposium on Theory of computing*, 2011.

[6] Chandra Chekuri and Alina Ene. Approximation algorithms for submodular multiway partition. In *FOCS*, 2011.

[7] Chandra Chekuri, Jan Vondrak, and Rico Zenklusen. Dependent randomized rounding via exchange properties of combinatorial structures. In *2010 IEEE 51st Annual Symposium on Foundations of Computer Science*, pp. 575–584. IEEE, 2010.

[8] Andrew Cotter, Mahdi Milani Fard, Seungil You, Maya Gupta, and Jeff Bilmes. Constrained interacting submodular groupings. In *International Conference on Machine Learning*, pp. 1068–1077. PMLR, 2018.

[9] J. Edmonds. Submodular functions, matroids and certain polyhedra. *Combinatorial structures and their Applications*, 1970.

[10] M.L. Fisher, G.L. Nemhauser, and L.A. Wolsey. An analysis of approximations for maximizing submodular set functions—II. In *Polyhedral combinatorics*, 1978.

[11] Mohammad Ghodsi, MohammadTaghi HajiAghayi, Masoud Seddighin, Saeed Seddighin, and Hadi Yami. Fair allocation of indivisible goods: Improvements and generalizations. In *Proceedings of the 2018 ACM Conference on Economics and Computation*, pp. 539–556, 2018.

[12] Daniel Golovin. Max-min fair allocation of indivisible goods. *Technical Report CMU-CS-05-144*, 2005.

[13] S. Jegelka and J. Bilmes. Submodularity beyond submodular energies: coupling edges in graph cuts. In *CVPR*, 2011.

[14] Subhash Khot and Ashok Ponnuswami. Approximation algorithms for the max-min allocation problem. In *APPROX*, 2007.

[15] Diederik Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.

[16] A. Krause and C. Guestrin. A note on the budgeted maximization of submodular functions, 2005. URL `http://citeseerx.ist.psu.edu/viewdoc/summary?doi=10.1.1.130.3314`.

[17] A. Krause, A. Singh, and C. Guestrin. Near-optimal sensor placements in Gaussian processes: Theory, efficient algorithms and empirical studies. In *JMLR*, 2008.

[18] Andreas Krause, Brendan McMahan, Carlos Guestrin, and Anupam Gupta. Robust submodular observation selection. In *JMLR*, 2008.

[19] Alex Krizhevsky and Geoffrey Hinton. Learning multiple layers of features from tiny images. Technical report, University of Toronto, 2009.

[20] Chandrashekhar Lavania and Jeff Bilmes. Auto-summarization: A step towards unsupervised learning of a submodular mixture. In *Proceedings of the 2019 SIAM International Conference on Data Mining*, pp. 396–404. SIAM, 2019.

[21] Kiyohito Nagano, Yoshinobu Kawahara, and Satoru Iwata. Minimum average cost clustering. In *Advances in Neural Information Processing Systems*, pp. 1759–1767, 2010.

[22] Z. Svitkina and L. Fleischer. Submodular approximation: Sampling-based algorithms and lower bounds. In *FOCS*, 2008.

[23] Ivor W Tsang, James T Kwok, and Pak-Ming Cheung. Core vector machines: Fast svm training on very large data sets. In *JMLR*, 2005.

[24] Jan Vondrák. Optimal approximation for the submodular welfare problem in the value oracle model. In *STOC*, 2008.

[25] Shengjie Wang, Wenruo Bai, Chandrashekhar Lavania, and Jeff Bilmes. Fixing mini-batch sequences with hierarchical robust partitioning. In *The 22nd International Conference on Artificial Intelligence and Statistics*, pp. 3352–3361, 2019.

[26] Kai Wei, Yuzong Liu, Katrin Kirchhoff, and Jeff Bilmes. Using document summarization techniques for speech data subset selection. In *North American Chapter of the Association for Computational Linguistics/Human Language Technology Conference (NAACL/HLT-2013)*, Atlanta, GA, June 2013.

[27] Kai Wei, Yuzong Liu, Katrin Kirchhoff, Chris Bartels, and Jeff Bilmes. Submodular subset selection for large-scale speech training data. In *Proc. IEEE Intl. Conf. on Acoustics, Speech, and Signal Processing*, Florence, Italy, 2014.

[28] Kai Wei, Rishabh Iyer, and Jeff Bilmes. Submodularity in data subset selection and active learning. In *ICML*, 2015.

[29] Kai Wei, Rishabh K Iyer, Shengjie Wang, Wenruo Bai, and Jeff A Bilmes. Mixed robust/average submodular partitioning: Fast algorithms, guarantees, and applications. In *Advances in Neural Information Processing Systems*, pp. 2233–2241, 2015.

[30] Jingjing Zheng, Zhuolin Jiang, Rama Chellappa, and Jonathon P Phillips. Submodular attribute selection for action recognition in video. In *NIPS*, 2014.

8 Neurips 2021 Paper Checklist

1. Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?

   Answer: yes.

2. Have you read the ethics review guidelines and ensured that your paper conforms to them?

   Answer: yes.

3. Did you discuss any potential negative societal impacts of your work?

   Answer: N/A. The main contribution of the paper is on the theoretical aspect. There might be indirect societal impacts due to the underlining machine learning applications, and we believe they are all benevolent.

4. Did you describe the limitations of your work?

   Answer: Yes. We discuss the trade-offs between the computation complexity and theoretical guarantees for the presented algorithms. We also discuss in Section 3.1 about the limitations on objects/bounds compared to the heterogeneous setting.

5. Did you state the full set of assumptions of all theoretical results?

   Answer: Yes. They are all stated clearly in Section 3 and stated in the theorems/lemmas.

6. Did you include complete proofs of all theoretical results?

   Answer: Yes. The complete proofs can be found in the appendix.

7. Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)?

   Answer: details about the experiments are described in Appendix

8. Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)?

   Answer: Yes.

9. Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)?

   Answer: Yes.

10. Did you include the amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)?

    Answer: Yes.

11. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...

    Answer: N/A.

12. If you used crowdsourcing or conducted research with human subjects... Answer: N/A.

We restate the theoretical statements and the algorithms here for completeness and convenience.

A Proofs for Minimum Block Greedy Based Algorithm

Given a normalized monotone submodular function $f : 2^V \to \mathbb{R}_{\geq 0}$, where V is the ground set with $|V| = n$ elements, and a positive integer m denoting the number of blocks in the partition, the Submodular Robust partition is defined as:

$$\max_{\pi \in \Pi(V,m)} \min_{A \in \pi} f(A), \tag{1}$$

Where $\Pi(V, m)$ is the set of all possible partitions on set V into m blocks, and π is one partition. The objective aims to find the partition such that the minimum-valued block in the partition is maximized. Denote the optimal partition to Eq. (1) as $\pi^* = \{O_1, O_2, \ldots, O_m\}$.

Lemma 2 (Unconstrained Min-Block Streaming[29]). *For a ground set V and its elements $(v_1, v_2, \ldots, v_n)$ coming in an arbitrary streaming order, the output solution of Alg. 1 has $\min_{j\in[m]} f(A_j) \geq \frac{1}{m} \min_{j\in[m]} f(O_j)$.*

Corollary 1 (Unconstrained Min-Block Greedy[29]). *The output solution of Alg. 2 has $\min_{j\in[m]} f(A_j) \geq \frac{1}{m} \min_{j\in[m]} f(O_j)$ since the order of adding elements in Min-Block Greedy is one possible order of the ground set elements.*

Lemma 3 (Tightness of Corollary 1). *$\forall \epsilon > 0$, $\exists$ a submodular function f such that the output solution of Alg 2 $\min_{j=1:m} f(A_j) = \frac{1}{m} \min_{j=1:m} f(O_j) + \epsilon$.*

Proof. We construct a set cover function as the tight example for Corollary 1. We illustrate the set cover function graphically in Fig. 3.

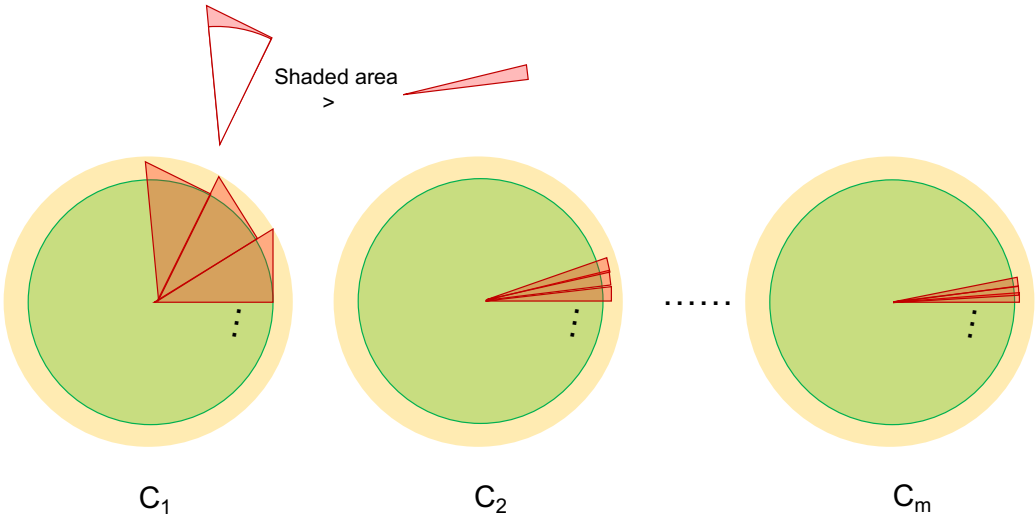

Figure 3: A graphical illustration of the tight example. The circles are the areas to cover for the set cover function and the green inner circles and the red triangles are elements in the ground set (the outer yellow circles are not elements). The inner circles (green) largely overlap with the outer circles (yellow). The red triangles mostly overlap with the inner circle, with little gains on the ring between the two circles. We can change the size of the red triangles so that Min-Block Greedy prefers a redundant element (the shaded area comparison on the top of the figure). Also note that the red triangles may overlap on the inner circle part (they may not retain the shapes as triangles), so overall they cover $m - 1$ times the area of each circle.

Suppose we have m circles of area to cover in the set cover function. Say we order the circles by their area, say $C_1 < C_2 < \ldots < C_m$. Let $C_{j+1} = C_j + \epsilon_{j+1}$ for some $\epsilon_{j+1} > 0$. For every circle j, we have an element $v_j \in V$, which covers an inner circle, which almost covers the entire circle. W.l.o.g, suppose $f(v_1) = 1$ and let $C_j = f(v_j) + \delta_j$.

For each circle C_j, we construct n_j elements, which largely overlap with the inner circle covered by v_j and gives little gain on the ring between the inner circle and the outer circle C_j. Call these n_j elements V_j. Let $f(V_j|v_j) = \epsilon'_j$, $f(v) < f(v_j) \forall v \in V_j$, and $f(V_j) > f(v_j)$.

Now let's focus on the first two circles C_1 and C_2, and assume $m = 2$ for the partition problem. It is easy to extend to general m case by recursively applying the following arguments on C_2 and C_3.

Suppose we run the min-block greedy, after the first two steps, one block contains v_1 and the other contains v_2. At step 3, v_1 is the min-block. By setting the suitable values for n_1 and n_2 (say $n_2 >> n_1$), we can make $f(v|v_1) > f(v'|v_1) \forall v \in V_1, v' \in V_2$. Therefore, we will still select an element from V_1 even though such element overlaps largely with the inner circle v_1. We can force the min-block greedy algorithm to select all elements from V_1 before the min-block changes to the block containing v_2, and $f(V_1 \cup v_1) > f(v_2)$. After that, the algorithm can only add elements from V_2 to the block containing v_2, which gives only ϵ'_2 gains. As we can make the values of ϵ_j, ϵ'_j and δ_j arbitrarily small, the solution is arbitrarily close to 1. On the contrary, the optimal partition should add elements in V_1 to v_2 and elements in V_2 to v_1, and the solution has value arbitrarily close to 2.

To extend to general m partitions, we may treat the current V_2 as V_1, and construct V_3 in the same way we construct V_2 based on V_1. In the first m steps of the min-block greedy, the algorithm is forced to evenly distribute $v_j, j = 1, 2, \ldots, m$ into every block. After that, the algorithm adds all elements in V_j to the block containing v_j before the min-block changes and the block containing v_j will not become the min-block again. In the end, every block only covers (almost) one circle. Suppose for V_j, we may make the elements to cover the inner circle multiple times, i.e., $\exists \pi \in \Pi(V_j, m) s.t. \forall A \in \pi, f(A) \geq f(v_j)$. Then for the optimal solution, every block can cover (almost) all the circles, and therefore the approximation ratio can be arbitrarily close to $1/m$.

□

The constrained submodular robust partition problem:

$$\max_{\pi \in \Pi(V,m,\mathcal{C})} \min_{A \in \pi} f(A), \tag{4}$$

Where $\Pi(V, m, \mathcal{C})$ is the set of all possible partitions on set V into m blocks such that for every partition $\pi \in \Pi(V, m, \mathcal{C})$, every block $A \in \pi$ should satisfy the constraint $A \in \mathcal{C}$. Denote the optimal partition to Eq. (4) as $\pi^*_C = \{O^{\mathcal{C}}_1, O^{\mathcal{C}}_2, \ldots, O^{\mathcal{C}}_m\}$.

Theorem 1 (Constrained Min-block Greedy). *Given a constraint $\mathcal{C}$, if the greedy solution S^g to problem $\max_{S \in \mathcal{C}} f(S)$ using Alg. 4 has a bound of α, i.e., $f(S^g) \geq \alpha \max_{S \in \mathcal{C}} f(S)$, then the solution of Alg. 3 has $\min_{j \in [m]} f(A_j) \geq \frac{\alpha}{\alpha m + 1} \min_{j \in [m]} f(O^{\mathcal{C}}_j)$. Assuming GreedyStep$(\cdot)$ takes $\mathcal{O}(1)$ oracle calls, the time complexity of of Alg. 4 is $\mathcal{O}(n^2)$.*

Proof. W.l.o.g., we assume that the block of index 1 for a partition corresponds to the minimum-valued block, e.g., $f(O^{\mathcal{C}}_1) = OPT^{\mathcal{C}}$. For Min-Block Greedy algorithm, we always add an element feasible to the constraint $\mathcal{C}$ to the block with the minimum evaluation. Let the minimum block in our final solution be A_1. Due to the final singleton comparison step (line 9-11 in Alg. 3), there are several different scenarios for A_1:

1. It is never the case that we cannot add any elements to a block due to the constraint (line 5 always true). This is the simplest case as we can directly reduce it to a stream of elements with the same ordering as we add them into different blocks, and Lemma. 2 applies. We therefore can get an $1/m$ approximation ratio, which is better than the one given in the theorem for any $\alpha \leq 1$.

2. A_1 is the first block that we cannot find any feasible elements to add. The singleton comparison step may increase the function value of A_1. however, by assumption it's still the minimum block after the algorithm completes.

3. There are other blocks that we cannot find any feasible elements to add before A_1. This could only happen if the other blocks get their values increased by the singleton comparison step. As if the singleton comparison step does not swap the block with the largest singleton, the block, which is not A_1 in this case, is the minimum block for that step and remains minimum for the following steps of the algorithm.

For scenarios 2 and 3, the general idea of the proof is the same, where we separate the ground set V into two parts V' and R' ($V = V' \cup R'$), and bound $f(A_1)$ by comparing to a block in the optimal solution $O_j^{\mathcal{C}}$ through $f(O_j^{\mathcal{C}} \cap V')$ and $f(O_j^{\mathcal{C}} \cap R')$. However, for 2 and 3, we will use slightly different V' and R'.

First, for scenario 2), let's suppose at step t', the current minimum block is A_1, and we find no feasible elements to add. Let all the elements allocated so far (before the singleton comparison step for A_1) as V', and the remaining unallocated elements as R'. $V = V' \cup R'$. Denote the elements in A_1 before the singleton comparison step as A_1', as the singleton step always improves the block value, we have $f(A_1) \geq f(A_1')$.

If we run the min-block robust partition greedy algorithm on V' only, we will get the same partial partition as we run on V for t' steps. Therefore, suppose we create a stream that orders the elements in V' in the same order that those elements get allocated by the min-block robust partition greedy algorithm, then by Lemma. 2, we have:

$$f(A_1) \geq f(A_1') \geq \frac{1}{m} OPT(V'), \tag{7}$$

Where we denote $OPT(V') = \max_{\pi \in \Pi(V',m)} \min_{A \in \pi} f(A)$ as the optimal solution for the unconstrained robust submodular partition on the ground set V'.

Let $O_j^{\mathcal{C}}$ be some block in the optimal constrained partition on ground set V. since $O_j^{\mathcal{C}}$ can be the non-minimal block in the optimal solution, we have:

$$f(O_j^{\mathcal{C}}) \geq OPT^{\mathcal{C}}. \tag{8}$$

There exists a $j \in \{1, \ldots, m\}$ such that

$$f(A_1) \geq \frac{1}{m} OPT(V') \tag{9}$$

$$\geq \frac{1}{m} f(O_j^{\mathcal{C}} \cap V'), \tag{10}$$

as otherwise $\forall j \in \{1, \ldots, m\}, O_j^{\mathcal{C}} \cap V'$ forms a solution for the partition problem on the reduced ground set V', and gives a solution value better than $OPT(V')$, which violates the optimality of $OPT(V')$.

Now we separate the constrained optimal solution on ground set V into 2 parts: $O_j^{\mathcal{C}} \cap V'$ and $O_j^{\mathcal{C}} \cap R'$.

Assumption 1. Suppose

$$f(O_j^{\mathcal{C}} \cap R') \geq f(O_j^{\mathcal{C}} \cap V'), \tag{11}$$

Then because of submodularity, $f(O_j^{\mathcal{C}} \cap R') + f(O_j^{\mathcal{C}} \cap V') \geq f(O_j^{\mathcal{C}})$ (recall $V = V' \cup R'$)and we have

$$f(O_j^{\mathcal{C}} \cap R') \geq \frac{1}{2} f(O_j^{\mathcal{C}}). \tag{12}$$

Consider the set $R' \cup A_1'$, let

$$\hat{O} \in \operatorname*{argmax}_{S \subseteq R' \cup A_1', S \in \mathcal{C}} f(S), \tag{13}$$

I.e., $\hat{O}$ is the optimal solution to the constraint submodular max on the reduced ground set $R' \cup A_1'$. After the singleton comparison step on A_1', we get A_1, which is the greedy solution of Alg. 3 on the reduced ground set $R' \cup A_1'$ and constraint $\mathcal{C}$. Therefore, based on the α-bound assumption in Theorem 1, we have:

$$f(A_1) \geq \alpha f(\hat{O}) \tag{14}$$

$$\geq \alpha f(O_j^{\mathcal{C}} \cap R') \tag{15}$$

$$\geq \frac{\alpha}{2} f(O_j^{\mathcal{C}}) \tag{16}$$

$$\geq \frac{\alpha}{2} f(O_1^{\mathcal{C}}). \tag{17}$$

Eq. (15) comes from the optimality of $\hat{O}$ and Eq. (16) comes from **Assumption 1**.

Assumption 2. Otherwise, we have

$$f(O_j^{\mathcal{C}} \cap V') > f(O_j^{\mathcal{C}} \cap R') \tag{18}$$

$$\geq \frac{1}{2} f(O_j^{\mathcal{C}}). \tag{19}$$

We therefore have:

$$f(A_1) \geq \frac{1}{m} f(O_j^{\mathcal{C}} \cap V') \tag{20}$$

$$> \frac{1}{2m} f(O_j^{\mathcal{C}}) \tag{21}$$

$$\geq \frac{1}{2m} f(O_1^{\mathcal{C}}) \tag{22}$$

Note that one of **Assumption 1** and **Assumption 2** is always true, since $f(O_j^{\mathcal{C}} \cap R') + f(O_j^{\mathcal{C}} \cap V') \geq f(O_j^{\mathcal{C}})$ because of submodularity. Previously, we use equal weights of $\frac{1}{2}$ for both assumptions. We can balance the weights as long as the weights sum to one, and we get:

if $f(O_j^{\mathcal{C}} \cap R') \geq \frac{1}{\alpha m+1} f(O_j^{\mathcal{C}})$, we have

$$f(A_1) \geq \alpha f(\hat{O}) \tag{23}$$

$$\geq \alpha f(O_j^{\mathcal{C}} \cap R') \tag{24}$$

$$\geq \frac{\alpha}{\alpha m + 1} f(O_j^{\mathcal{C}}) \tag{25}$$

$$\geq \frac{\alpha}{\alpha m + 1} f(O_1^{\mathcal{C}}); \tag{26}$$

if $f(O_j^{\mathcal{C}} \cap V') > \frac{\alpha m}{\alpha m+1} f(O_j^{\mathcal{C}})$, we have

$$f(A_1) > \frac{1}{m} f(O_j^{\mathcal{C}} \cap V') \tag{27}$$

$$> \frac{1}{m} \frac{\alpha m}{\alpha m + 1} f(O_j^{\mathcal{C}} \cap V) \tag{28}$$

$$> \frac{\alpha}{\alpha m + 1} f(O_1^{\mathcal{C}}). \tag{29}$$

Thus, we get a $\frac{\alpha}{\alpha m+1}$ bound.

For scenario 3, we only need to change V' and R' and the same argument follows. Recall that in such a scenario, there are some other blocks that have no feasible elements to add before A_1, and they get their values increased through the singleton comparison step. There are also two different cases here. Firstly, it does not happen that there are no feasible elements to add to A_1 until the end of the algorithm. In such a case, similar to the scenario 1, we can order the elements as a stream and applies Lemma. 2 to get the $1/m$ approximation ratio. Note that the blocks that get to the singleton comparison step all get their values increased for scenario 3, and we can just add the singleton a^* (line 9) to that block in the streaming case. To be more precise, for Alg. 3 when block j ($j \neq 1$) is the current minimum block, and has no feasible elements to add, we denote its elements as A'_j, and the singleton comparison step gives an element a^* with $f(\{a^*\}) \geq f(A'_j)$. In the streaming ordering, we use the same ordering as we add element in Alg. 3, and at the singleton comparison step for block j, we have the next element in the stream be a^*, and we add that element to block j since block j is the current minimum block. By monotonicity, we have $f(A'_j \cup \{a^*\}) \geq f(\{a^*\}) \geq f(A_1)$. In other words, those blocks never become the minimum block again, and no elements get added to them after their singleton comparison step. Therefore, we have a streaming ordering of the elements that will make the minimum block equal to A_1 and Lemma. 2 applies.

Next, we discuss for the case where it happens that there are no feasible elements to add to A_1. When that happens, we set all the allocated elements as V' and the remaining elements as R' before the singleton comparison step. Note for those blocks that get to the singleton comparison step before A_1, we will also include the singletons in V' Recall those singletons have larger gains and get swapped

with the elements in those blocks for this scenario. As stated above, those blocks with the singleton comparison step never become the minimum block again, so they don't interfere with the remaining blocks/elements. For such V' and R', the exact argument in scenario 2 can be made, i.e., we can treat A'_1 as a min-block streaming solution on V' (Eq. 20) and A_1 as a greedy solution on $A'_1 \cup R'$ (Eq. 15).

□

Corollary 2 (**Cardinality Constrained Min-block Greedy).** *For $\mathcal{C}$ as a cardinality constraint, the output of Alg 3 has* $\min_{j=1:m} f(A_j) \geq \frac{1}{m+\frac{1}{1-e^{-1}}} \min_{j=1:m} f(O_j^{\mathcal{C}})$.

Corollary 3 (**Matroid Constrained Min-block Greedy).** *For $\mathcal{C}$ as an intersection of p matroids constraint, the output of Alg 3 has* $\min_{j=1:m} f(A_j) \geq \frac{1}{m+p+1} \min_{j=1:m} f(O_j^{\mathcal{C}})$.

Corollary 4 (**Knapsack Constrained Min-block Greedy).** *For $\mathcal{C}$ as a knapsack constraint, the output of Alg 3 has* $\min_{j=1:m} f(A_j) \geq \frac{1}{m+\frac{2}{1-1/e}} \min_{j=1:m} f(O_j^{\mathcal{C}})$.

B Proofs for Round-Robin Greedy Based Algorithms

The matroid constrained submodular robust partition problem is

$$\max_{\pi \in \Pi(V,m,\mathcal{M})} \min_{A \in \pi} f(A). \tag{30}$$

Before we get into the proofs for the algorithm bounds, we will state the following lemma, which is a general property about robust submodular partitioning.

Lemma 5 (**Removal of one element and one block).** *For any $v \in V$, we have:*

$$\max_{\pi \in \Pi(V \setminus v,m-1,\mathcal{M})} \min_{A \in \pi} f(A) \geq \max_{\pi \in \Pi(V,m,\mathcal{M})} \min_{A \in \pi} f(A). \tag{31}$$

I.e., if we remove one element and one block from the problem, the optimal solution gets no worse.

Proof. Denote the optimal solution on V and m by $O_1, O_2, \ldots, O_m$ with $f(O_1) \leq f(O_2) \leq \ldots \leq f(O_m)$.

Suppose $v \in O_j$ for some j, then the blocks other than O_j forms a solution for problem defined on $V \setminus v$ and $m-1$, and we can add elements in $O_j \setminus v$ to other blocks (if the constraints permit). In the worst case, even if we cannot add any elements of $O_j \setminus v$ to other blocks, we still have $\max_{\pi \in \Pi(V,m,\mathcal{M})} \min_{A \in \pi} f(A) \geq \min_{j' \in [m], j' \neq j} f(O_{j'}) \geq f(O_1)$.

Suppose $\forall j \in [m],\ v \notin O_j$, then we only remove one block, and we can add the elements in that block to any other block so the solution value gets improved.

□

For Round-Robin Greedy based algorithms, we first guess the optimal solution value and then assign singletons to blocks which satisfies the bound based on the guessed optimal value. After that, we run the algorithm on the restricted problem with those blocks and elements removed. By applying the previous lemma (recursively if multiple elements and blocks removed), we know that the optimal solution on the restricted instance is no worse than the optimal solution on the original problem. Therefore, it suffices to analyze the solution on the restricted instance.

Lemma 4 (**Cardinality Constrained Round-Robin).** *For the problem in Eq.* (4)*, with $\mathcal{C}$ as a cardinality constraint k, Alg. 5 gives a solution* $\min_{j \in [m]} f(A_j) \geq \frac{(1-e^{-1})^2}{3} \min_{j \in [m]} f(O_j^k)$.

Proof. By solving $\max_{\pi \in \Pi(V',m',k)} \sum_{S \in \pi} f(S)$ in Line 10 of Alg. 5 (Theorem III.3 in [7]), we know that

$$\sum_{j \in [m]'} f(A'_j) \geq (1-e^{-1}) \max_{\pi \in \Pi(V',m',k)} \sum_{S \in \pi} f(S). \tag{32}$$

Algorithm 5: Cardinality Round-Robin Greedy

input : f, V, m, cardinality constraint k, discounting factor for guessing optimal δ

Let τ be the solution value of Alg. 3;

Let $high = \lceil \log_{1+\delta}(m+2) \rceil$, $low = 0$;

Create a sequence of guessed values: $(\tau, (1+\delta)\tau, (1+\delta)^2\tau, \ldots, (1+\delta)^{high}\tau)$;

Create an empty solution ($\emptyset$ for each block in the partition) for each guessed value $\pi_0, \pi_1, \ldots, \pi_{high}$;

while $high \geq low$ **do**

  Let $idx = \lfloor (high + low)/2 \rfloor$; Let $A_1 = A_2 = \ldots = A_m = \emptyset$;

  Let $V' = \{v | v \in V, f(v) \leq \frac{(1-e^{-1})^2}{3}(1+\delta)^{idx}\tau\}$; Let $G = V \setminus V'$;

  Assign G to $A_{m-|G|+1}, A_{m-|G|+2}, \ldots, A_m$ with one element per block;

  Let $m' = m - |G|$;

  Let $A'_1, A'_2, \ldots, A'_{m'}$ be the solution to $\max_{\pi \in \Pi(V', m', k)} \sum_{S \in \pi} f(S)$ using continuous greedy and swap rounding; Let $V'' = \cup_{j \in [m]'} A'_j$;

  Let $\{A_1, A_2, \ldots, A_{m'}\} = RR(f, V'', m', \mathcal{M}_k^u, [m'])$;

  if $f(A_j) \geq \frac{(1-e^{-1})^2}{3}(1+\delta)^{idx}\tau \; \forall j \in [m']$ **then**

    Let $\pi_{idx} = \{A_1, A_2, \ldots, A_m\}$; Let $low = idx + 1$;

  else

    Let $high = idx - 1$;

return best of $\pi_0, \pi_1, \ldots, \pi_{high}$;

Recall that we denote the optimal solution value in the cardinality constraint case by $OPT^{\mathcal{M}_k^u}$, where k is the cardinality. We assume we know the optimal solution value $OPT^{\mathcal{M}_k^u}$ for this proof. For the algorithm, the $OPT^{\mathcal{M}_k^u}$ value is guessed within a factor of $\frac{1}{1+\delta}$. Therefore, to be more precise, we have an additional factor of $\frac{1}{1+\delta}$ in the bound, which can be made arbitrarily small by setting δ small.

For the limited ground set V'', running unconstrained round-robin ensures every block to have at most k elements and therefore the cardinality constraint is satisfied. Suppose we run the continuous greedy algorithm on the limited ground set V'' with the submodular welfare objective ($\max_{\pi \in \Pi(V'', m')} \sum_{S \in \pi} f(S)$), we get a fractional solution $x_1 = x_2 = \ldots = x_{m'} = (\frac{1}{m'}, \frac{1}{m'}, \ldots, \frac{1}{m'})$ (we do not really need to run the algorithm, but we will compare our solution to the fractional solution). Denote the multilinear extension of f by F and $F(x) = \mathbb{E}_{R \sim x} f(R)$ (we can think it as the expected value of f where every element is sampled independently based on probabilities defined in vector x). Consider any block A_j in the solution of Alg 5 for $j \in [m']$ (for $j \notin [m']$, those blocks are the singleton assignment blocks and they satisfy the bound by construction), we have:

$$\frac{(1-e^{-1})^2}{3} OPT^{\mathcal{M}_k^u} + 2f(A_j) \geq F(x_j) \tag{33}$$

$$\geq \frac{1-e^{-1}}{m'} \max_{\pi \in \Pi(V'', m')} \sum_{S \in \pi} f(S) \tag{34}$$

$$\geq \frac{1-e^{-1}}{m'} \max_{\pi \in \Pi(V'', m', k)} \sum_{S \in \pi} f(S) \tag{35}$$

$$\geq \frac{(1-e^{-1})^2}{m'} \max_{\pi \in \Pi(V', m', k)} \sum_{S \in \pi} f(S) \tag{36}$$

$$\geq (1-e^{-1})^2 \max_{\pi \in \Pi(V', m', k)} \min_{S \in \pi} f(S). \tag{37}$$

$$\geq (1-e^{-1})^2 OPT^{\mathcal{M}_k^u}. \tag{38}$$

Rearrange and we get:

$$f(A_j) \geq \frac{(1-e^{-1})^2}{3} OPT^{\mathcal{M}_k^u}. \tag{39}$$

Eq. (33) comes from the Lemma.3 of [3], in which case we can bound every block in the round-robin solution to the fractional solution of the continuous greedy algorithm on the multilinear extension of f. Note for Lemma.3 of [3], they study the unconstrained case and show that $\frac{\gamma}{3}OPT^{\mathcal{M}_k^u} + 2f(A_j) \geq F(x_j)$ with $\gamma = 1 - e^{-1}$. γ comes from the singleton assignment step, which assigns blocks with singletons whose values are larger than $\frac{\gamma}{3}OPT^{\mathcal{M}_k^u}$. A slightly more general statement can be made for any $0 \leq \gamma \leq 1$ with the same proof as Lemma.3 of [3]. In our case, we pick $\gamma = (1 - e^{-1})^2$. Eq. (34) follows from the property of the continuous greedy solution. The continuous greedy gives a $(1 - e^{-1})$ approximation to the submodular welfare problem, and the fractional solution x_j for each block is the same. Therefore, every block's evaluation in expectation is at least $\frac{(1-e^{-1})}{m'}$ of the submodular welfare optimal solution. Eq. (35) follows that the unconstrained solution is no worse than the constrained solution. Eq. (36) uses Eq. (32): $A'_1, \ldots, A'_m$ is one possible solution to $\max_{\pi \in \Pi(V'', m', \mathcal{M}_k^u)} \sum_{S \in \pi} f(S)$, and we know $\sum_j f(A'_j) \leq \max_{\pi \in \Pi(V'', m', k)} \sum_{S \in \pi} f(S)$ because of the max operator. Therefore, we have $\max_{\pi \in \Pi(V'', m', k)} \sum_{S \in \pi} f(S) \geq \sum_j f(A'_j) \geq (1 - e^{-1}) \max_{\pi \in \Pi(V', m', k)} \sum_{S \in \pi} f(S)$. Eq. (37) follows that the sum over blocks of the max-min solution is no larger than the optimal welfare solution. Eq. (38) uses Lemma. 5: the optimal max-min solution on V' and m' is no worse than the optimal max-min solution on V and m.

As stated above, $\frac{\gamma}{3}OPT + 2f(A_j) \geq F(x_j)$ is true for any $0 \leq \gamma \leq 1$. It may seem that making γ smaller can improve the bound. However, we only bound the m' blocks but not the singleton blocks. Because of the singleton assignment step, every singleton has a value larger than $\frac{\gamma}{3}$, and the final bound over all m blocks will be the minimal of the bound on the m' blocks and the singleton blocks. Setting γ small worsens the bound on the singleton blocks. To balance the two bounds, $\gamma = (1 - e^{-1})^2$ is picked so that the bounds on the m' blocks and the singleton blocks meet.

Finally, we discuss about the binary search process. Let's consider a special problem instance, which does not have any singleton values larger than $\frac{(1-1/e)^2}{3}$ times the optimal solution value of the problem instance. Then we can directly run line 10 and 11 to get a solution with a $\frac{(1-1/e)^2}{3}$ approximation ratio (no optimal solution value guessing and large singleton value removal). However, we don't know if that assumption is true for general problem instances.

Back to the general problem instances. Suppose the guessed optimal values form a sequence $(\tau_1, \tau_2, \ldots, \tau_l)$ where $\tau_{i+1} = (1 + \delta)\tau_i$. We are going to show that for any $\tau_i \leq OPT^{\mathcal{M}_k^u}$ as the guessed optimal value that we plug into the binary search iterations (Alg. 5 line 7-15), line 12 of Alg. 5 is always true. For simplicity, denote the optimal solution of the cardinality constrained robust partitioning problem as OPT for the following. Let's denote the found large singleton values and the remaining sets (line 7 of Alg. 5) respectively by G_{OPT} and V'_{OPT} for OPT, and G_i and V'_i for τ_i. Since $\tau_i \leq OPT$, $G_{OPT} \subseteq G_i$ as the threshold is smaller for τ_i. Then by Lemma. 5, we know that the optimal solution OPT_i on V'_i (partitioned to $m - |G_i|$ blocks) is no less than the optimal solution OPT' on V'_{OPT} (partitioned to $m - |G_{OPT}|$ blocks). Also note that since we remove all singleton values to form V'_i based on the threshold $\frac{(1-e^{-1})^2}{3}\tau_i$, and $\tau_i \leq OPT \leq OPT' \leq OPT_i$ ($OPT \leq OPT'$ is also from Lemma. 5), we are guaranteed that there are no singleton elements with $f(v) \geq \frac{(1-e^{-1})^2}{3}OPT_i$ in V'_i. Therefore, as discussed for the special problem instance, line 10 and 11 on V'_i give a solution whose every block has a value of at least $\frac{(1-e^{-1})^2}{3}OPT_i \geq \frac{(1-e^{-1})^2}{3}\tau_i$, and thus line 12 is guaranteed to be true for τ_i.

Based on that, either of the following must hold: 1) for all $\tau_i > OPT^{\mathcal{M}}$ line 12 is false, and in such a case, we can use binary search to find the largest τ_i with line 12 true, and let's call it τ_{i^*}; 2) there exists some $\tau_i > OPT^{\mathcal{M}}$ such that line 12 is true. If we find such τ_i, we find a solution with $f(A_j) \geq \frac{(1-e^{-1})^2}{3}\tau_i \geq \frac{(1-e^{-1})^2}{3}OPT^{\mathcal{M}_k^u}$. Otherwise if we don't find it, we go to the first case and will still find τ_{i^*}. We can therefore conclude that binary search can be applied to search among the guessed optimal solution values.

□

B.1 Proof for the Matroid Constraint Case

Lemma 6 (**Continuous Greedy Solution**). *For the constrained welfare problem* $\max_{\pi \in \Pi(V,m,\mathcal{M})} \sum_{A \in \pi} f(A)$*, the continuous greedy algorithm outputs a fractional solution* $x_1 = x_2 = \ldots = x_m$ *(*$x_j \in [0,1]^n$*), which is the same for every block in the partition and* $\sum_{j \in [m]} F(x_j) \geq (1 - e^{-1}) \max_{\pi \in \Pi(V,m,\mathcal{M})} \sum_{A \in \pi} f(A)$*.* F *is the multilinear extension of* f*, i.e.,* $F(x) = \mathbb{E}_{R \sim x} f(R)$ *(we can think it as the expected value of* f *where every element is sampled independently based on probabilities defined in vector* x*). Moreover,* $\forall i \in [n], j \in \{1, \ldots, m\}, x_j[i] \leq \frac{1}{m}$ *and* $\sum_{i \in V} x_j[i] \leq r_{\mathcal{M}}(V)$*.*

Proof. Note that the continuous greedy can give a fractional solution with $1 - e^{-1}$ bound under any solvable polytope constraint. It's the rounding procedure that limits the constraint we can use to get a set solution, e.g., with pipage rounding, we can use any matroid constraint.

In fact, we do not need to run the continuous greedy algorithm, and we only need to show the existence of a solution. Suppose the solution to the $max_y\{w \cdot y, y \in \mathcal{P}\}$ step of the continuous greedy algorithm is given by some oracle. Given the direction y, we just evenly split the resulting vector y among the m blocks, as we cannot distinguish between blocks. At the end of the algorithm, we will have the fractional solution $x_1 = x_2 = \ldots = x_m$ and $\sum_{j \in [m]} F(x_j) \geq (1 - e^{-1}) \max_{\pi \in \Pi(V,m,\mathcal{M})} \sum_{A \in \pi} f(A)$.

Since the fractional solution are guaranteed to be in the matroid polytope of $\mathcal{M}$ and $\mathcal{M}_m^p$, we have $\forall i \in [n], j \in \{1, \ldots, m\}, x_j[i] \leq \frac{1}{m}$ and $\sum_{i \in V} x_j[i] \leq r_{\mathcal{M}}(V)$.

□

Lemma 7 (**Matroid Constraint Round-robin with No Large Singletons**). *Suppose for all* $v \in V$*, we have* $f(v) \leq \frac{1-e^{-1}}{5} OPT^{\mathcal{M}}$*, where* $OPT^{\mathcal{M}}$ *is the optimal solution value of the robust submodular partition problem constrained by matroid* $\mathcal{M}$ *(in other words, all the singletons have relatively small values for the given problem instance). Then, the round-robin iterations of Alg. 7 (line 10) gives a solution* $\min_{j \in [m]} f(A_j) \geq \frac{1-e^{-1}}{5} OPT^{\mathcal{M}}$*.*

Algorithm 7: Matroid Round-Robin Greedy

input : f, V, m, matroid constraint $\mathcal{M}$, discounting factor for guessing optimal δ

1 Let τ be the solution value of Alg. 3;
2 Let $high = \lceil \log_{1+\delta}(m+2) \rceil$, $low = 0$;
3 Create a sequence of guessed values: $(\tau, (1+\delta)\tau, (1+\delta)^2\tau, \ldots, (1+\delta)^{high}\tau)$;
4 Create an empty solution ($\emptyset$ for each block in the partition) for each guessed value $\pi_0, \pi_1, \ldots, \pi_{high}$;
5 **while** $high \geq low$ **do**
6 Let $idx = \lfloor (high + low)/2 \rfloor$; Let $A_1 = A_2 = \ldots = A_m = \emptyset$;
7 Let $V' = \{v | v \in V, f(v) \leq \frac{1-e^{-1}}{5}(1+\delta)^{idx}\tau\}$; Let $G = V \setminus V'$;
8 Assign G to $A_{m-|G|+1}, A_{m-|G|+2}, \ldots, A_m$ with one element per block;
9 Let $m' = m - |G|$;
10 Let $\{A_1, A_2, \ldots, A_{m'}\} = RR(f, V', m', \mathcal{M}, [m'])$;
11 **if** $f(A_j) \geq \frac{(1-e^{-1})}{5}(1+\delta)^{idx}\tau \; \forall j \in [m']$ **then**
12 Let $\pi_{idx} = \{A_1, A_2, \ldots, A_m\}$; Let $low = idx + 1$;
13 **else**
14 Let $high = idx - 1$;
15 **return** best of $\pi_0, \pi_1, \ldots, \pi_{high}$;

Proof. Let's focus on one block (any one in $A_1, \ldots, A_{m'}$) and for simplicity, we will omit the block index j for this proof if not further noticed. Denote $OPT = \min_{j \in [m]} f(O_j^{\mathcal{M}})$ for this proof. Also, we assume we know the optimal solution value OPT for this proof. Note that in the complete version of Alg. 7, we need to remove large singleton values based on the guessed optimal value, but for this lemma we make the assumption that in the given problem instance, there are no large singletons present.

For the current block, we denote the final resulting set from Alg. 7 as A. For one round-robin iteration, we go over all the feasible blocks sequentially, and to get A we need to run $|A| = r$ round-robin iterations. Note that for different blocks, the number of round-robin iterations might be different.

We then divide the restricted ground set V' by the round-robin iterations with respect to the current block A. Before we add the first element to A, denote all the allocated elements by V^0. Then we can think that for every round-robin iteration, we always start from the current block A. Let $V' = V^0 \cup V^1 \cup \ldots \cup V^r$ be a partition of V' and V^t contains all the elements allocated during the t's round-robin iteration. Note V^r contains all the unallocated elements in the ground set after we add the last element to A. Let $V^{t_1:t_2} = \cup_{t\in\{t_1,t_1+1,\ldots,t_2\}} V^t$. Accordingly, we partition the result A by $A^t = A \cap V^t$.

For the set $V' \setminus A$, we separate it into two parts Q'_1 and Q'_2, where Q'_1 contain all the elements checked in Alg. 7 that cannot be added to the current block due to the matroid constraint, and let $Q'_2 = V' \setminus A \setminus Q'_1$, i.e., Q'_2 contain all the elements that can be added the current block. To be more precise:

$$Q'_1 = \cup_{t\in\{0,1,\ldots,r\}} \cup_{v\in V^t\setminus A^t,(A^{1:t}\cup v)\notin\mathcal{M}} v \tag{40}$$

Let $Q_1 = Q'_1 \cup A$ and $Q_2 = Q'_2 \cup A$.

Let F denote the multilinear extension of f, i.e., $F(x) = \mathbb{E}_{R\sim x} f(R)$. By Lemma 5 and Lemma 6, we know that

$$\begin{aligned}
(1-e^{-1})OPT &\leq \max_{\pi\in\Pi(V',m',\mathcal{M})} \min_{S\in\pi} f(S) && (41)\\
&\leq \frac{1}{m'} \max_{\pi\in\Pi(V',m',\mathcal{M})} \sum_{S\in\pi} f(S) && (42)\\
&\leq F(x) && (43)\\
&\leq (F(x\cap Q_1) + F(x\cap Q_2)). && (44)
\end{aligned}$$

Note that x is the fractional solution to the continuous greedy algorithm on the welfare objective: $\max_{\pi\in\Pi(V',m',\mathcal{M})} \sum_{S\in\pi} f(S)$ (similar to one of the x_j's in Lemma 6 and we omit the block index for this proof). Here we use $x \cap Q$ to represent setting all elements not in Q to be zero in the x fractional solution. The first inequality follows from Lemma. 5. The second inequality follows that the sum over blocks of the max-min solution is no better than the optimal solution of the welfare problem. Since every element is sampled independently according to its probability in the fractional solution x, together with submodularity ($Q_1 \cup Q_2 = V'$) we get the last inequality above. Next, we will bound the two terms $F(x \cap Q_1)$ and $F(x \cap Q_2)$ separately.

For the first term $F(x \cap Q_1)$, we know that $r = r_{\mathcal{M}}(Q_1)$, and Alg. 7 generates A in the same manor as running greedy max on Q_1 with matroid constraint $\mathcal{M}$. To be more precise, suppose $\mathcal{M} = (V, \mathcal{I})$ and we remove all the elements that are not in Q_1 and get $\mathcal{M}' = (V' \cap Q_1, \{I \cap Q_1 \forall I \in \mathcal{I}\})$. Note that $\mathcal{M}'$ is also a matroid, and all sets that satisfy $\mathcal{M}'$ also satisfy $\mathcal{M}$ due to the down-monotone property of matroids. Therefore, we have:

$$\begin{aligned}
f(A) &\geq \frac{1}{2} \max_{S\in\mathcal{M}'} f(S) && (45)\\
&\geq \frac{1}{2} F(x\cap Q_1). && (46)
\end{aligned}$$

Note that x is in the matroid polytope of $\mathcal{M}$, and $x \cap Q_1$ is in the matroid polytope of $\mathcal{M}'$. By pipage rounding, we know that we can get an integral solution X' from $F(x \cap Q_1)$ so that the integral solution still satisfies $X' \in \mathcal{M}'$ and $f(X') \geq F(x \cap Q_1)$. Since $\max_{S\in\mathcal{M}'} f(S) \geq f(X')$, we get the last inequality above.

For the second term $F(x \cap Q_2)$, we will bound it using the greedy step. Denote $y = x \cap Q_2$, $y^t = y \cap V^t$, and $\mathbb{E}_S(y) = \mathbb{E}_{R\sim y} f(R|S)$, we have:

$$
\begin{aligned}
F(y) &= \mathbb{E}_{R\sim y} f(R) && (47)\\
&= \mathbb{E}_{R\sim y^0} f(R) + \mathbb{E}_{R^1\sim y^0, R^2\sim y^{1:r}} f(R^2|R^1) && (48)\\
&\leq F(y^0) + F(y^{1:r}) && (49)\\
&\leq F(y^0) + f(A^1) + \mathbb{E}_{A^1}(y^{1:r}) && (50)\\
&= F(y^0) + f(A^1) + \mathbb{E}_{A^1}(y^1) + \mathbb{E}_{R^1\sim y^1, R^2\sim y^{2:r}} f(R^2|R^1 \cup A^1) && (51)\\
&\leq F(y^0) + f(A^1) + \mathbb{E}_{A^1}(y^1) + \mathbb{E}_{A^1}(y^{2:r}) && (52)\\
&\leq F(y^0) + f(A^1) + \mathbb{E}_{A^1}(y^1) + f(A_2|A_1) + \mathbb{E}_{A^2}(y^{2:r}) && (53)
\end{aligned}
$$

Continue to unwrap $\mathbb{E}_{A^2}(y^{2:r})$ in the same way, finally we get:

$$
\begin{aligned}
F(y) \leq\ & F(y^0) + \left[f(A^1) + f(A^2|A^1) + f(A^3|A^2) + \ldots + f(A^r|A^{r-1})\right]\\
& + \left[\mathbb{E}_{A^1}(y^1) + \mathbb{E}_{A^2}(y^2) + \ldots + \mathbb{E}_{A^r}(y^r)\right] && (54)\\
=\ & F(y^0) + f(A) + \left[\mathbb{E}_{A^1}(y^1) + \mathbb{E}_{A^2}(y^2) + \ldots + \mathbb{E}_{A^r}(y^r)\right] && (55)
\end{aligned}
$$

We then need to bound $F(y^0)$ and $\left[\mathbb{E}_{A^1}(y^1) + \mathbb{E}_{A^2}(y^2) + \ldots + \mathbb{E}_{A^r}(y^r)\right]$. Note that by assumption, we do not have any singleton gains larger than $\frac{1-e^{-1}}{5}$, and we have:

$$
F(y^0) \leq \frac{1-e^{-1}}{5} OPT \tag{56}
$$

Since we select items greedily at every round-robin step, and y only has non-zero values for elements that are in Q_2, we have:

$$
\begin{aligned}
\mathbb{E}_{A^t}(y^{t+1}) &= \mathbb{E}_{R\sim y^{t+1}} f(R|A^t) && (57)\\
&\leq \sum_{v\in y^{t+1}} y^{t+1}(v) f(v|A^t) && (58)\\
&\leq \sum_{v\in y^{t+1}} \frac{1}{m'} f(v|A^t) && (59)\\
&\leq f(A^t|A^{t-1}) && (60)
\end{aligned}
$$

Note that for the last round-robin iteration V^r, it may seem that there can be more than m' elements, but it's not possible: since there are no new elements added to the current blocks, $V^r \cap Q_2$ contains at most m' elements as otherwise we will find new feasible elements and add to block A.

Then we sum over all t and get:

$$
\left[\mathbb{E}_{A^1}(y^1) + \mathbb{E}_{A^2}(y^2) + \ldots + \mathbb{E}_{A^r}(y^r)\right] \leq f(A). \tag{61}
$$

Therefore, we have:

$$
\begin{aligned}
(1-e^{-1})OPT &\leq (F(x\cap Q_1) + F(x \cap Q_2)) && (62)\\
&\leq 2f(A) + F(y) && (63)\\
&\leq 2f(A) + 2f(A) + \frac{1-e^{-1}}{5} OPT && (64)\\
f(A) &\geq \frac{1-e^{-1}}{5} OPT. && (65)
\end{aligned}
$$

□

Next, we will discuss why binary search can be used in the guessing of the optimal value (as opposed to the case of linear search where we try all possible guessed optimal values). We make almost the same argument as the cardinality constraint case. We restate the argument and proof here for completeness.

Lemma 8 (**Binary Search**). *For Alg. 7, let the potential guessed optimal values form a sequence* $(\tau_1, \tau_2, \ldots, \tau_l)$ *where* $\tau_{i+1} = (1+\delta)\tau_i$. *Then for any* $\tau_i \leq OPT^{\mathcal{M}}$ *as the guessed optimal value that we plug into the binary search iterations (Alg. 7 line 7-14), line 14 of Alg. 7 is always true.*

Proof. For simplicity, denote the optimal solution of the matroid constrained robust partitioning problem as OPT for this proof. Let's denote the found large singleton values and the remaining sets (line 7 of Alg. 7) respectively by G_{OPT} and V'_{OPT} for OPT, and G_i and V'_i for τ_i. Since $\tau_i \leq OPT$, $G_{OPT} \subseteq G_i$ as the threshold is smaller for τ_i. Then by Lemma. 5, we know that the optimal solution OPT_i on V'_i (partitioned to $m - |G_i|$ blocks) is no less than the optimal solution OPT' on V'_{OPT} (partitioned to $m - |G_{OPT}|$ blocks). Also note that since we remove all singleton values to form V'_i based on the threshold $\frac{1-e^{-1}}{5}\tau_i$, and $\tau_i \leq OPT \leq OPT' \leq OPT_i$ ($OPT \leq OPT'$ is also from Lemma. 5), we are guaranteed that there are no singleton elements with $f(v) \geq \frac{1-e^{-1}}{5}OPT_i$ in V'_i. Therefore, based on Lemma. 7, our round robin iterations on V'_i give a solution whose every block has a value of at least $\frac{1-e^{-1}}{5}OPT_i \geq \frac{1-e^{-1}}{5}\tau_i$, and thus line 11 is guaranteed to be true for τ_i. □

Based on Lemma. 8, either of the following must hold: 1) for all $\tau_i > OPT^{\mathcal{M}}$ line 11 is false, and in such a case, we can use binary search to find the largest τ_i with line 11 true, and let's call it τ_{i^*}; 2) there exists some $\tau_i > OPT^{\mathcal{M}}$ such that line 11 is true. If we find such τ_i, we find a solution with $f(A_j) \geq \frac{1-e^{-1}}{5}\tau_i \geq \frac{1-e^{-1}}{5}OPT^{\mathcal{M}}$. Otherwise if we don't find it, we go to the first case and will still find τ_{i^*}.

Finally, the approximation guarantee of Alg.7 follows by combining the previous lemmas.

Theorem 2 (**Matroid Constrained Round-Robin**). *For the problem in Eq.* (4)*, with* $\mathcal{C}$ *as any matroid constraint* $\mathcal{M}$*, Alg 7 gives a solution* $\min_{j\in[m]} f(A_j) \geq \frac{(1-e^{-1})}{5}\min_{j\in[m]} f(O_j^{\mathcal{M}})$. *The time complexity of of Alg. 7 is* $\mathcal{O}(n^2(\log\log m + \log\frac{1}{\delta}))$.

Proof. Firstly, based on the previous arguments about the binary search, we can find a τ_i with $\tau_i \geq \tau_{i^*}$, where τ_{i^*} is the largest τ_j with $\tau_j \leq OPT^{\mathcal{M}}$ and line 11 of Alg. 7 true. By setting δ small, τ_{i^*} can be arbitrarily close to $OPT^{\mathcal{M}}$. Next, based on Lemma. 7, after removing the large singleton values, the solution on the remaining elements have the min block value at least $\frac{1-e^{-1}}{5}\tau_i$, and the removed large singletons are all larger than $\frac{1-e^{-1}}{5}\tau_i$ by construction. We therefore get the approximation ratio. □

C Synthetic Experiment

We compare Alg. 3, Alg. 7 (10 optimal value guesses) and the random assignment baseline on randomly generated synthetic facility location functions. Every entry in the facility location similarity matrix is uniformly sampled from $[0, 1]$. We report our results with different parameters in Figure 4.

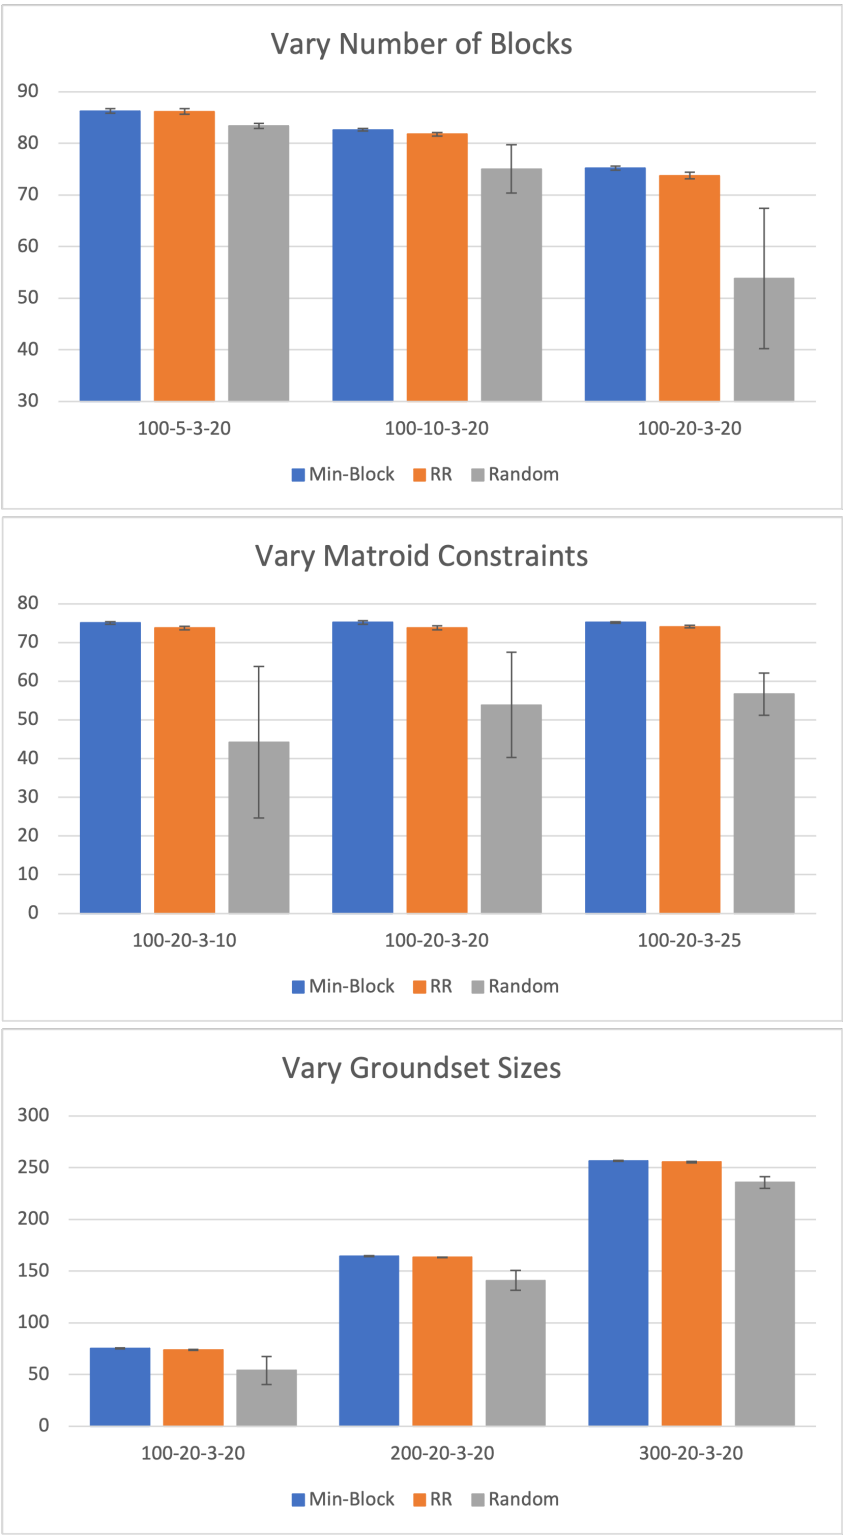

Figure 4: Synthetic data results on randomly generated facility location function similarity matrices. The $x-$label is $n - m - c - p$. Results averaged over 30 runs.

Particularly, we have four parameters for every problem instance: n the ground set size, m the number of partitions, and the partition matroid constraint parameters c and p. For the partition

matroid constraint, we divide the groundset into blocks of size p and for every such block, it is constrained that we can pick at most c elements. For all variants of the settings, we observe that Alg. 3 and Alg. 7 significantly outperform the random baseline. Alg. 3 consistently outperforms Alg. 7 but the margin is not very large.

D Experiment Details

The features used in the experiments are generated through an autoencoder. The network architecture is described in Table 1. The network is trained using ReLU non-linearity and batch normalization. ADAM [15] is utilized as the optimization method with an initial learning rate of 5e-3, a weight decay of 5e-4 and a minibatch size of 100. The network is trained in PyTorch using the procedure described in [20]. Features are extracted as the output of the autoencoder's bottleneck (following the residual block and non-linearity).

The training of the ResNet-9 (Myrtle-AI, `https://github.com/davidcpage/cifar10-fast`) network utilizes an ADAM optimizer with an initial learning rate of $1e-3$. The network is trained for 90 epochs. For CPU jobs, we use a single core Intel Xeon 2.10GHz CPU, and for GPU jobs we use a Nvidia RTX 2080Ti GPU.

Table 1: Neural network structure of the autoencoder

Group	Block Type (kernel sz, stride, channels)	# Blocks
conv1	$[\,3\times 3\,]$, 2, 64	1
conv1 (residual)	$\begin{bmatrix} 3\times 3 \\ 3\times 3 \end{bmatrix}$, 1, 64	2
conv2	$[\,3\times 3\,]$, 2, 16	1
conv2 (residual)	$\begin{bmatrix} 3\times 3 \\ 3\times 3 \end{bmatrix}$, 1, 16	2
conv3	$[\,3\times 3\,]$, 2, 8	1
conv3 (residual)	$\begin{bmatrix} 3\times 3 \\ 3\times 3 \end{bmatrix}$, 1, 8	2
conv4	$[\,3\times 3\,]$, 1, 4	1
conv4 (residual)	$\begin{bmatrix} 3\times 3 \\ 3\times 3 \end{bmatrix}$, 1, 4	1
deconv4 (residual)	$\begin{bmatrix} 3\times 3 \\ 3\times 3 \end{bmatrix}$, 1, 4	1
deconv3	$[\,3\times 3\,]$, 1, 8	1
deconv3 (residual)	$\begin{bmatrix} 3\times 3 \\ 3\times 3 \end{bmatrix}$, 1, 8	2
deconv2	$[\,3\times 3\,]$, 2, 16	1
deconv2 (residual)	$\begin{bmatrix} 3\times 3 \\ 3\times 3 \end{bmatrix}$, 2, 16	2
deconv1	$[\,3\times 3\,]$, 2, 64	1
deconv1 (residual)	$\begin{bmatrix} 3\times 3 \\ 3\times 3 \end{bmatrix}$, 2, 64	2
deconv0	$[\,3\times 3\,]$, 2, 3	1

We also test the case for various cardinality constraints on CIFAR-10 datatset, and report the results in Figure 5. The running times (in seconds) for Figure 5 and Figure 1 is given in Table 2 and Table 3 respectively under various constraint parameters. For the round-robin algorithm, we pick $\delta = 0.1$. The computing platform uses a single core of Intel Xeon(R) CPU 2.10GHz.

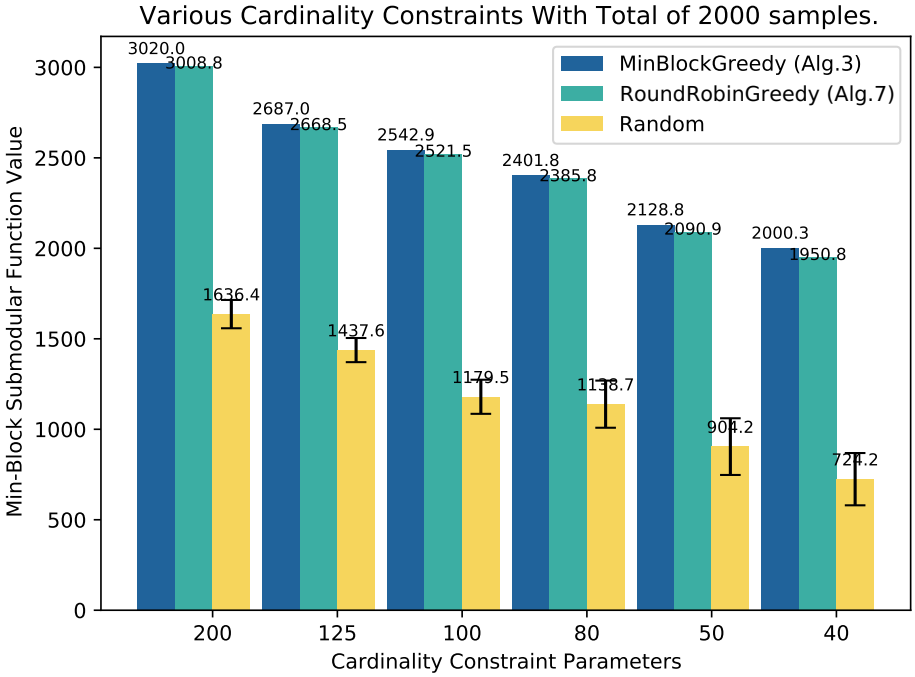

Figure 5: Cardinality constrained results. We select a total of 2000 samples and the number of blocks varies according to the constraint, e.g., when the constraint $k = 200$, $m = 2000/200 = 10$.

Table 2: Wall Clock Running Time for Figure 5

Constraint Param	200	125	100	80	50	40
Min-Block Greedy	60.7	60.4	60.0	57.9	57.0	57.1
Round-Robin Greedy	240.2	233.5	243.7	244.6	240.0	242.1

Table 3: Wall Clock Running Time for Figure 1

Constraint Param	40	25	20	10	8	5
Min-Block Greedy	63.4	62.3	62.2	60.0	59.5	59.9
Round-Robin Greedy	184.2	180.3	182.1	242.4	244.0	240.0