
Near Optimal Algorithms for Hard Submodular Programs with Discounted Cooperative Costs

Rishabh Iyer
Microsoft Corporation

Jeff Bilmes
University of Washington, Seattle

Abstract

In this paper, we investigate a class of submodular problems which in general are very hard. These include minimizing a submodular cost function under combinatorial constraints, which include cuts, matchings, paths, etc., optimizing a submodular function under submodular cover and submodular knapsack constraints, and minimizing a ratio of submodular functions. All these problems appear in several real world problems but have hardness factors of $\Omega(\sqrt{n})$ for general submodular cost functions. We show how we can achieve constant approximation factors when we restrict the cost functions to low rank sums of concave over modular functions. A wide variety of machine learning applications are very naturally modeled via this subclass of submodular functions. Our work therefore provides a tighter connection between theory and practice by enabling theoretically satisfying guarantees for a rich class of expressible, natural, and useful submodular cost models. We empirically demonstrate the utility of our models on real world problems of cooperative image matching and sensor placement with cooperative costs.

1 Introduction

Submodular functions provide a rich class of expressible models for a variety of machine learning problems. Submodular functions occur naturally in two flavors. In minimization problems, they model notions of cooperation, attractive potentials, and economies of scale, while in maximization problems, they model aspects of coverage, diversity, and information. A set function $f : 2^V \rightarrow \mathbb{R}$ over a finite set $V = \{1, 2, \dots, n\}$ is *submodular* [4] if for all subsets $S, T \subseteq V$, it holds that

Proceedings of the 22nd International Conference on Artificial Intelligence and Statistics (AISTATS) 2019, Naha, Okinawa, Japan. PMLR: Volume 89. Copyright 2019 by the author(s).

$f(S) + f(T) \geq f(S \cup T) + f(S \cap T)$. Given a set $S \subseteq V$, we define the *gain* of an element $j \notin S$ in the context S as $f(j|S) \triangleq f(S \cup j) - f(S)$. A perhaps more intuitive characterization of submodularity is as follows: a function f is submodular if it satisfies *diminishing marginal returns*, namely $f(j|S) \geq f(j|T)$ for all $S \subseteq T, j \notin T$, and is *monotone* if $f(j|S) \geq 0$ for all $j \notin S, S \subseteq V$.

In this paper, we address the following a family of hard submodular optimization problems. The first one is constrained submodular minimization [6, 9, 13, 15, 16, 29]:

$$\text{Problem 1: } \min\{f(X) | X \in \mathcal{C}\}$$

where the function f is *monotone submodular*, and \mathcal{C} is a *combinatorial constraint*, which could represent a cardinality lower bound constraint, or more complicated ones like cuts, matchings, trees, or paths in a graph. With cut constraints, this problem becomes cooperative cuts [17], and with matching constraints, we call this *cooperative matchings*, which we introduce and utilize in this paper.

The second problem asks for minimizing a monotone submodular cost function f , while simultaneously maximizing a monotone submodular coverage function g . A natural way to model this bi-optimization problem is to introduce one of f and g as a constraint [12]. In particular, we obtain two optimization problems:

$$\text{Problem 2: } \min\{f(X) | g(X) \geq c\},$$

$$\text{Problem 3: } \max\{g(X) | f(X) \leq b\},$$

The fourth problem considered in this paper is minimizing the ratio of submodular functions [1].

$$\text{Problem 4: } \min\{f(X)/g(X) | \emptyset \subset X \subset V\}$$

A key assumption in this paper is that the functions f and g in Problems 1-4, are monotone submodular – an assumption that, as we shall see, is natural in many applications. Problem 2 is a special case of Problem 1, with $\mathcal{C} = \{X : g(X) \geq c\}$. Furthermore, Problem 2 and Problem 3 are closely related and, loosely speaking, duals of each other [12]. Similarly, Problem 4



Figure 1: An illustration of cooperative matchings. The goal is to match corresponding points between the two images. On the left, bipartite matching is used and there are many mismatched points (indicated by red edges). On the right, similar points are clustered and offered a within-cluster discount via a submodular function, significantly reducing the number of mis-matched points.

is closely related to Problems 2 and 3, in that given an approximation algorithm for Problems 2 or 3, we can obtain an approximation with similar guarantees for Problem 4 [1] (also considered in [27] with general monotone set functions). Problem 1 is constrained submodular minimization, while Problems 2, 3 and 4 try to simultaneously minimize one submodular function while maximizing another.

Problems 1-4 appear naturally in several machine learning applications. However, in the worst case all four problems have polynomial hardness factors of $\Omega(\sqrt{n})$ [29, 15, 12, 1]. An important observation is that the polynomial hardness of problems 1 - 3, comes up mainly due to the submodular cost function f – they do not depend as much on the constraints \mathcal{C} or the submodular function g [15, 12]. In the case of Problem 4, the hardness depends on both f and g [1].

On the other hand, these problems come up as models in many machine learning applications. These lower bounds are specific to rather contrived classes of functions, whereas much better results can be achieved for many practically relevant cases. The pessimistic worst case results are somewhat discouraging, begging the need to quantify sub-classes of submodular functions that are more amenable to these optimization problems. Only limited past work has focused on investigating these problems with potentially a subclass of submodular functions. [14, 12, 1] provide bounds for Problems 1-4 based on the notion of *curvature*, and argue how several submodular functions (e.g. clustered concave over modular functions) have bounded curvature. Their curvature bounds depend on the choice of the submodular functions, and in certain cases yield no improvement over the worst case bounds. For classes of functions with bounded curvature, their bounds yield improved results.

In this paper, we focus on a tractable yet expressive

subclass of submodular cost functions f , namely low rank sums of concave over modular functions.

Definition 1. *Low rank sums of concave over modular functions are the class of functions representable as $f(X) = \sum_{i=1}^k \psi_i(w_i(X))$, where ψ_i s are monotone concave, and k is constant or $O(\log n)$.*

Low Rank in this context means that the number of components in the sum is small (i.e., k is small). Our use of the terminology “low rank” is identical to that used in [8]. We argue how this subclass naturally models many interesting applications of problems 1 - 4 in machine learning. We do not need to consider the entire class of submodular functions (which includes rather contrived instances), but only this subclass. This observation helps us in providing better connections between theory and practice. The main specialty of this subclass is that these functions effectively model cooperation between objects via discounts provided by concave functions. Moreover, we show that this subclass admits fully polynomial time approximation schemes for Problem 1, and constant factor approximation guarantees for Problems 2 and 3. Similarly, we achieve constant factor approximation guarantees for Problem 4, when f is a low rank sum of concave over Modular functions, and g is an arbitrary submodular function, a significant improvement over [1]. The bounds we obtain are significantly better than the worst case bounds, and also an improvement over the bounds achieved using the curvature [15, 12].

Low rank sums of concave over modular functions in Problems 1 - 4, fit as natural models in several machine learning problems. Below, we summarize some of these.

Image segmentation (Cooperative Cuts): Markov random fields with pairwise attractive potentials occur naturally in modeling image segmentation and related applications [2]. While models are tractably solved using graph-cuts, they suffer from the shrinking bias problem, and images with elongated edges are not segmented properly. When modeled via a submodular function, however, the cost of a cut is not just the sum of the edge weights, but a richer function that allows cooperation between edges, and yields superior results on many challenging tasks (see, for example, the results of the image segmentations in [17]). This was achieved in [17] by partitioning the set of edges \mathcal{E} of the grid graph into groups of similar edges (or types) $\mathcal{E}_1, \dots, \mathcal{E}_k$, and defining a function $f(S) = \sum_{i=1}^k \psi_i(w(S \cap \mathcal{E}_i))$, $S \subseteq \mathcal{E}$, where ψ_i s are concave functions and w encodes the edge potentials. This ensures that we offer a *discount* to edges of the same type. Moreover, the number of types of edges are typically much smaller than the number of pixels, so this is a low-rank sum of concave functions.

Image Correspondence (Cooperative Match-

ings): The simplest model for matching key-points in pairs of images (which is also called the correspondence problem) can be posed as a bipartite matching. These models, however, do not capture interaction between the pixels. We illustrate the difficulty of this in Figure 1. One kind of desirable interaction is that similar or neighboring pixels be matched together. We can achieve this as follows. First we cluster the key-points in the two images into k groups (this is illustrated in Figure 1-left via green, blue and red key-points). This induces a clustering of edges that can be given a discount via a submodular function (details are given in Section 4.1). In practice, the number of groups (k) can be much smaller than n and this is a low-rank sum of concave over modular functions. Figure 1-right shows how the submodular matchings improves over the simple bipartite matching. In particular, the minimum matching approach produces many spurious matches between clusters (shown in red) that are avoided via the cooperation described above.

Sensor Placement or Feature Selection: Often, the problem of choosing sensor locations A from a given set of possible locations V can be modeled [22, 10] by maximizing the mutual information between the chosen variables A and the unchosen set $V \setminus A$ (i.e., $g(A) = I(X_A; X_{V \setminus A})$). Alternatively, we may wish to maximize the mutual information between a set of chosen sensors X_A and a quantity of interest C (i.e., $g(A) = I(X_A; C)$) assuming that the set of features X_A are conditionally independent given C [22]. Both these functions are submodular. Since there are costs involved, we want to simultaneously minimize the cost $f(A)$. Often this cost is submodular [22, 10], since there is typically a discount when purchasing sensors in bulk (or computing features), and we can express this via Problems 2 and 3. For example, there may be diminished cost for placing a sensor in a particular location given placement in certain other locations. Similarly, certain features might be cheaper to use given that others are already being computed (e.g., those that use an FFT). A natural cost model in such cases is $f(A) = \sum_{i=1}^k \psi_i(m(A \cap S_i))$ where ψ_i 's are concave, $m(j)$ is the cost of sensor (or feature) j and S_1, \dots, S_k are groups of similar sensors or features. Typically, k is much smaller than n and this can be expressed as low rank sum of concave over modular functions.

2 Background & Existing Algorithms

The basic idea for most combinatorial algorithms solving Problems 1 - 4, are based on approximating the cost function f with a tractable surrogate function \hat{f} [6, 7, 16, 13, 12, 15, 1]. Moreover, all four problems have similar guarantees. We characterize the quality of the solution via the notion of approximation factors. In particu-

lar, we say that an algorithm achieves an approximation factor of $\alpha \geq 1$ for Problem 1, if we can obtain a set \hat{X} such that $f(\hat{X}) \leq \alpha f(X^*)$, where X^* is the optimizer of Problem 1. For Problems 2 and 3, we use the notion of bi-criterion approximation factors. An algorithm is a $[\sigma, \rho]$ bi-criterion algorithm for Problem 2 if it is guaranteed to obtain a set \hat{X} such that $f(\hat{X}) \leq \sigma f(X^*)$ (approximate optimality) and $g(\hat{X}) \geq \rho c$ (approximate feasibility), where X^* is an optimizer of Problem 2. Typically, $\sigma \geq 1$ and $\rho \leq 1$. Similarly, an algorithm is a $[\rho, \sigma]$ bi-criterion algorithm for Problem 3 if it is guaranteed to obtain a set \hat{X} such that $g(\hat{X}) \geq \rho g(X^*)$ and $f(\hat{X}) \leq \sigma b$, where X^* is the optimizer of Problem 3. Moreover, problems 2 and 3 are very closely related [12], in that an approximation algorithm for one problem can be used to obtain guarantees for the other problem. The two problems also have matching hardness factors. For Problem 4, we study an algorithm which achieve α -approximation guarantees, in that we can achieve a set \hat{X} such that $h(\hat{X}) \leq \alpha h(X^*)$ where $h(X) = f(X)/g(X)$ and X^* is the optimal minimizer of h .

Supergradient based Algorithm (SGA): One such method uses the supergradients of a submodular function [15, 13, 6, 17, 11] to obtain modular upper bounds in an iterative manner. In particular, define a modular upper bound:

$$m_X^f(Y) \triangleq f(X) - \sum_{j \in X \setminus Y} f(j|V \setminus j) + \sum_{j \in Y \setminus X} f(j|X) \geq f(Y)$$

The algorithm starts with the $X^0 = \emptyset$ and sequentially sets X^{i+1} as the solution of the corresponding problem (1, 2 or 3) with a surrogate function as $\hat{f}(X) = m_{X^i}^f(X)$ [15, 16, 12]. In each case, this subproblem is much easier. For example, in the case of Problem 1, the subproblem becomes,

$$X^{i+1} = \min\{m_{X^i}^f(X) | X \in \mathcal{C}\}, \quad (1)$$

which is a linear cost problem, poly-time solvable for many constraints, like cardinality, cuts, matchings, paths etc.

In the case of Problems 2 and 3, these subproblems are

$$\begin{aligned} X^{i+1} &= \min\{m_{X^i}^f(X) | g(X) \geq c\} \text{ and} \\ X^{i+1} &= \max\{g(X) | m_{X^i}^f(X) \leq b\}, \end{aligned}$$

which are the submodular set cover and the submodular knapsack problems respectively [31, 24, 12], and are constant factor approximable to a factor of $1 - 1/e$.

With Problem 4, the subproblem becomes,

$$X^{i+1} = \min\{m_{X^i}^f(X)/g(X) | \emptyset \subset X \subset V\}$$

This can be approximated up to a factor of $e/(e-1)$ via a Greedy algorithm [1].

Lemma 1. Define $\alpha_f(X^*) = \frac{|X^*|}{1+(|X^*|-1)(1-\hat{\kappa}_f(X^*))} \leq \min\{|X^*|, \frac{1}{1-\hat{\kappa}_f(X^*)}\}$, where $\hat{\kappa}_f(X) = 1 - \frac{\sum_{j \in X} f(j|X \setminus j)}{\sum_{j \in X} f(j)}$ represents the average curvature of the function f . The supergradient based iterative algorithm (SGA) achieves an approximation factor of $\alpha_f(X^*)$ for Problem 1, and bicriteria factors satisfying $\sigma = \alpha_f(X^*)$ and $\rho = 1 - 1/e$ for Problems 2 and 3. Finally, SGA achieves an approximation factor of $e/(e-1) * \alpha_f(X^*)$ for Problem 4.

This Lemma follows easily from the results in [15, 13, 12, 1]. We can also achieve a non-bicriteria approximation factor for Problem 2, which is worse than the bicriteria factor by a log factor [12]. A key quantity which defines the approximation factor above is the average curvature $\hat{\kappa}_f(X^*)$, which in turn depends on the concave functions. If the concave function is $\psi_i(x) = x^a, a \in (0, 1)$, SGA admits approximation factors of $O(|X^*|^{1-a})$ [13]. On the other hand, if the concave function is $\psi(x) = \log(1+x)$, the guarantees are $O(|X^*|)$, which is much poorer.

The supergradient based algorithm is easy to implement, and also works well in practice [15, 17]. For the general class of submodular functions, these results are close to the optimal bounds, and are, in fact, tight for some constraints. Nevertheless, the worst case guarantees seem discouraging, particularly for the class of low rank sums of concave over modular functions that we consider here, and that as mentioned above are natural for many applications.

Ellipsoidal Approximation based Algorithm (EA): Another generic approximation of a submodular function, introduced by Goemans et. al [7], is based on approximating the submodular polyhedron by an ellipsoid. The main result states that any polymatroid (monotone submodular) function f , can be approximated by a function of the form $\sqrt{w^f(X)}$ for a certain modular weight vector $w^f \in \mathbb{R}^V$, such that $\sqrt{w^f(X)} \leq f(X) \leq O(\sqrt{n} \log n) \sqrt{w^f(X)}, \forall X \subseteq V$. A simple trick then provides a curvature-dependent approximation [13]. We have the following result borrowed from [12, 13, 1].

Lemma 2. Define $\alpha = O(\frac{\sqrt{n}}{1+(\sqrt{n}-1)(1-\kappa_f)})$, where $\kappa_f = 1 - \frac{\min_{j \in X} f(j|V \setminus j)}{\min_{j \in X} f(j)}$ represents the worst case curvature of the function f . The Ellipsoidal Approximation based algorithm (EA) achieves an approximation factor of α for Problem 1, and bicriteria factors satisfying $\sigma = \alpha$ and $\rho = 1 - 1/e$ for Problems 2 and 3. Similarly EA achieves an approximation guarantee of $e\alpha/(e-1)$ for Problem 4.

The Ellipsoidal Approximation obtains the tightest bounds for Problems 1-4 [12, 13, 7, 6, 16, 1]. This is again for the general class of submodular functions and the worst case factor of $O(\sqrt{n})$ is quite discouraging.

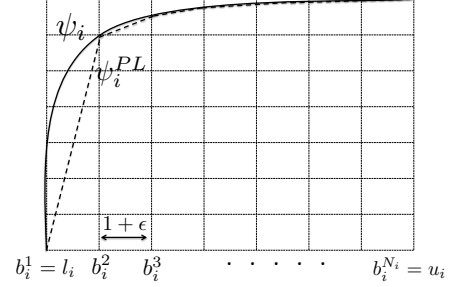


Figure 2: Visualizing ψ_i^{PL} and ψ_i .

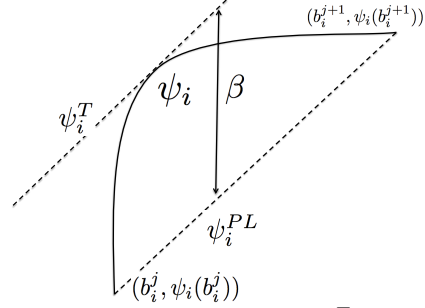


Figure 3: Showing ψ_i^T and ψ_i^{PL} .

This algorithm, however is very expensive computationally, and is not practical for solving machine learning applications [15].

3 Improved Algorithms for Low-rank sums of concave-modular functions

Our main new results are that we can achieve a fully polynomial time approximation scheme for Problem 1, and constant factor approximation guarantees for Problems 2, 3 and 4 when the cost function f is a low rank sum of concave over modular functions (Theorem 4). Our techniques build on recent methods used for minimizing quasi-concave functions over solvable polytopes [25, 23, 8, 19].

Assume the concave functions ψ_i 's are monotone functions, i.e., $\psi_i(y) \leq \psi_i(y'), \forall y \leq y'$. We also assume that for all i , $\psi_i(ky) \leq k^c \psi_i(y)$ for $k \geq 1, y \geq 0$ and some constant c . The second assumption holds for a number of concave functions, including $\psi_i(x) = x^a, a \in (0, 1)$, $\psi_i(x) = \log(1+x)$ and $\psi_i(x) = \min(x, a)$.

The main idea of this approach is to replace the concave functions ψ_i 's by piece-wise linear approximations $\psi_i^{PL}(x)$. We define an approximation of $f(X)$ as $f^{PL}(X)$ defined as $f^{PL}(X) = \sum_{i=1}^k \psi_i^{PL}(w_i(X))$. We then optimize this piece-wise linear approximation function, and the approximation factor comes based on the tightness of this piece-wise linear approximation. We call this procedure the *piece-wise linear approximation*

based algorithm (PLA).

We compute this approximation as follows. In the case of Problem 1, compute $l_i = \min\{w_i(X) | X \in \mathcal{C}\}$ and $u_i = \max\{w_i(X) | X \in \mathcal{C}\}$ for each $i = 1, 2, \dots, k$. Both these computations are linear cost problems and are polynomial time for most constraints. In case these are NP hard for Problem 1, or in the case of Problems 2, 3 and 4, we set $l_i = \min\{w_i(j), j \in V : w_i(j) > 0\}$ and $u_i = w_i(V)$. Then divide the range $[l_i, u_i]$ into pieces with breakpoints $b_i^1, b_i^2, \dots, b_i^{N_i}$ such that $b_i^1 = l_i$, $b_i^2 = l_i(1 + \epsilon)$, $b_i^3 = l_i(1 + \epsilon)^2$ and so on, for any $\epsilon > 0$. It is easy to see that $N_i = \log_{1+\epsilon} u_i/l_i \approx \log(u_i/l_i)/\epsilon$. The precision ϵ defines the fineness of the points, and the quality of the approximation.

For all $i = 1, 2, \dots, k$, define the piece-wise linear function ψ_i^{PL} , via the breakpoints $b_i^1, b_i^2, \dots, b_i^{N_i}$. A visualization of this is shown in Figure 2, where the dotted lines are the piece-wise approximation, while the solid curve is the concave function ψ_i . We first show that the function f^{PL} approximates the function f within a factor of $1 + \epsilon$.

Lemma 3. *The piece-wise linear function f^{PL} defined with a precision ϵ' satisfies,*

$$f^{PL}(X) \leq f(X) \leq (1 + \epsilon')^c f^{PL}(X) = (1 + \epsilon) f^{PL}(X) \quad (2)$$

where c is a constant such that $\psi_i(ky) \leq k^c \psi_i(y)$ for $k \geq 1, y \geq 0$ for all $i = 1, 2, \dots, k$.

Proof. By the construction of f^{PL} , and the concavity of the ψ_i s, it is easy to see that $f^{PL}(X) \leq f(X), \forall X \subseteq V$. To show the upper bound, consider a region defined by breakpoints b_i^j and b_i^{j+1} . Due to concavity of ψ_i , there exists a tangent at some point in $[b_i^j, b_i^{j+1}]$ whose slope equals that of the line connecting $(b_i^j, \psi_i(b_i^j))$ and $(b_i^{j+1}, \psi_i(b_i^{j+1}))$. This tangent line upper bounds the concave function ψ_i , and we can denote the corresponding upper bound as ψ_i^T . It then holds that $\psi_i^{PL}(y) \leq \psi_i(y) \leq \psi_i^T(y)$. We now show that $\psi_i^T(y) \leq (1 + \epsilon')^c \psi_i^{PL}(y)$.

We now focus on the region $[b_i^j, b_i^{j+1}]$. Let β be the constant difference between the two (parallel) lines, in terms of the y value. A visualization of this is shown in Figure 3. We would like to give a worst case bound on $\psi_i^T(y)/\psi_i^{PL}(y), \forall y \in [b_i^j, b_i^{j+1}]$. Notice that $\psi_i^T(y)/\psi_i^{PL}(y) = 1 + \beta/\psi_i^{PL}(y) \leq 1 + \beta/\psi_i^{PL}(b_i^j) = 1 + \beta/\psi_i(b_i^j) \leq \psi_i(b_i^{j+1})/\psi_i(b_i^j)$. The last inequality holds since $\psi_i(b_i^j) + \beta \leq \psi_i(b_i^{j+1})$, and the second last one holds since b_i^j is a break point.

Moreover, $b_i^{j+1} = b_i^j(1 + \epsilon')$ and hence $\psi_i(b_i^{j+1})/\psi_i(b_i^j) \leq \psi_i((1 + \epsilon')b_i^j)/\psi_i(b_i^j) \leq (1 + \epsilon')^c = 1 + \epsilon$. \square

We now show how we can exactly solve Problems 1, 2 and 3 using the cost function f^{PL} . Let s_i^j denote the slopes of the piece-wise linear functions – in other words, $s_i^j = [\psi_i(b_i^{j+1}) - \psi_i(b_i^j)]/[b_i^{j+1} - b_i^j]$. Also, we denote c_i^j as the corresponding intercepts. The functions ψ_i^{PL} are characterized by the pairs $\{(s_i^1, c_i^1), (s_i^2, c_i^2), \dots, (s_i^{N_i}, c_i^{N_i})\}$, and $\psi_i^{PL}(y) = s_i^j \cdot y + c_i^j, \forall y \in [b_i^j, b_i^{j+1}]$. We then consider the $\prod_{i=1}^k N_i$ different possibilities of the cross-terms. Define $J = [j_1, j_2, \dots, j_k]$ as a vector such that $J \in [1, N_1] \times [1, N_2] \times \dots \times [1, N_k]$.

In the case of Problem 1, PLA solves a set of optimization problems,

$$\hat{X}_J = \operatorname{argmin}\left\{\sum_{i=1}^k s_i^{j_i} w_i(X) + c_i^{j_i} \mid X \in \mathcal{C}\right\}, \quad \forall J \in [1, N_1] \times [1, N_2] \times \dots \times [1, N_k]. \quad (3)$$

The final solution \hat{X} is the minimum among the ones above. For problem 2, we consider the set of problems,

$$\hat{X}_J = \operatorname{argmin}\left\{\sum_{i=1}^k s_i^{j_i} w_i(X) + c_i^{j_i} \mid g(X) \geq c\right\}, \quad \forall J \in [1, N_1] \times [1, N_2] \times \dots \times [1, N_k], \quad (4)$$

and again set \hat{X} is the minimum among the \hat{X}_J 's above. Similarly, for Problem 3, we solve,

$$\hat{X}_J = \operatorname{argmax}\{g(X) \mid \sum_{i=1}^k s_i^{j_i} w_i(X) + c_i^{j_i} \leq b\}, \quad \forall J \in [1, N_1] \times [1, N_2] \times \dots \times [1, N_k] \quad (5)$$

We set \hat{X} corresponding to the set with the largest value of $g(\hat{X}_J)$. Finally, for Problem 4, we have:

$$\hat{X}_J = \operatorname{argmin}\left\{\frac{\sum_{i=1}^k s_i^{j_i} w_i(X) + c_i^{j_i}}{g(X)}\right\}, \quad \forall J \in [1, N_1] \times [1, N_2] \times \dots \times [1, N_k]. \quad (6)$$

Our main result is that these simple procedures provide improved guarantees for all three problems.

Theorem 4. *PLA achieves an approximation factor of $1 + \epsilon$ for Problem 1 as long as a linear function can be exactly minimized under \mathcal{C} . PLA also achieves a bi-criterion approximation factor satisfying $\sigma = 1 + \epsilon$ and $\rho = 1 - 1/e$ for Problems 2 and 3. PLA also achieves a non bicriterion approximation factor of $(1 + \epsilon) \log g(V)$ for Problem 2. PLA also achieves an approximation factor of $e(1 + \epsilon)/(e - 1)$ for Problem 4. The worst case complexity of PLA is $\prod_{i=1}^k \log(u_i/l_i)(\frac{1}{\epsilon})^k T = O((\frac{1}{\epsilon})^k T)$, where T is the complexity of Problems 1-4, with a linear cost function f .*

Proof. We first show that PLA solves Problems 1-4 with the surrogate function f^{PL} . Note that with the piece-wise linear approximation, Problem 1 becomes $\min_{X \in \mathcal{C}} \sum_{i=1}^k \psi_i^{PL}(w_i(X)) = \min_{X \in \mathcal{C}} \sum_{i=1}^k \min\{s_i^1 w_i(X) + c_i^1, s_i^2 w_i(X) + c_i^2, \dots, s_i^{N_i} w_i(X) + c_i^{N_i}\}$. This holds since $\psi_i^{PL}(y) = \min\{s_i^1 y + c_i^1, s_i^2 y + c_i^2, \dots, s_i^{N_i} y + c_i^{N_i}\}$, due to the concavity of ψ_i 's. We can then rewrite this as $\sum_{i=1}^k \psi_i^{PL}(w_i(X)) = \min_{J \in \mathcal{J}} \sum_{i=1}^k s_i^{j_i} w_i(X) + c_i^{j_i}$, where $\mathcal{J} = [1, N_1] \times [1, N_2] \times \dots \times [1, N_k]$. Combining these facts, we can rewrite the problem as $\min_{X \in \mathcal{C}} \min_{J \in \mathcal{J}} \sum_{i=1}^k s_i^{j_i} w_i(X) + c_i^{j_i}$, which after interchanging the min's becomes $\min_{J \in \mathcal{J}} \min_{X \in \mathcal{C}} \sum_{i=1}^k s_i^{j_i} w_i(X) + c_i^{j_i}$. This is exactly Eq. (3).

The algorithm for Problem 2 (Eq. (4)) is basically the same as that of Problem 1, since it is a special case. Similarly we can write Problem 4 as $\min_{X \in \mathcal{C}} \min_{J \in \mathcal{J}} \sum_{i=1}^k (s_i^{j_i} w_i(X) + c_i^{j_i})/g(X)$ which is equivalent to $\min_{J \in \mathcal{J}} \min_{X \in \mathcal{C}} \sum_{i=1}^k (s_i^{j_i} w_i(X) + c_i^{j_i})/g(X)$, which becomes equation (6). Equations (3), (4) and (5) each become instances of Problems 1, 2 and 4 with f being modular and the approximation guarantees follow directly from [14, 12, 1].

To deal with Problem 3, we use the fact that $f^{PL}(X) = \min_{J \in \mathcal{J}} \sum_{i=1}^k (s_i^{j_i} w_i(X) + c_i^{j_i})$, and hence we have the constraint, $\{\min_{J \in \mathcal{J}} \sum_{i=1}^k s_i^{j_i} w_i(X) + c_i^{j_i} \leq b\}$. First we show that X_J is feasible for all $J \in \mathcal{J}$. This follows easily from the fact that if for any X , $w_J(X) \leq b$, it holds that $\min_{J \in \mathcal{J}} w_J(X) \leq b$. Next, let X^* be the optimal solution of Problem 3, and let J^* be such that $\sum_{i=1}^k s_i^{j_i^*} w_i(X^*) + c_i^{j_i^*} = f^{PL}(X^*)$. Note that our algorithm covers J^* and hence $g(\hat{X}) \geq g(\hat{X}_{J^*}) \geq (1 - 1/e)g(X^*)$, where $1 - 1/e$ is the approximation factor of the submodular knapsack problem [28]. Note that the approximation factor of Problem 1 with f^{PL} is 1 assuming \mathcal{C} admits an exact solution with linear cost functions, while the factor for problem 2 is $\log g(V)$ for non-bicriterion algorithms, and a bicriterion factor of $[1, 1 - 1/e]$ with a bi-criterion algorithm [31, 12]. \square

Results similar to Theorem 4 have been shown for a generalization of Problem 1, which asks for constrained optimization of low rank functions [25, 23, 8, 19, 20, 5]. This problem in general is not a combinatorial optimization problem. However, when the functions are quasi-concave, the optimum lies on an extreme point, and hence, can be posed as a combinatorial optimization problem. Problem 1 asks for optimizing a specific subclass of concave (and hence quasi-concave) functions. [25, 8, 19] focus on the class of low rank quasi-concave functions, while [23] consider the general class of low rank functions. While their algorithms apply to our class of functions as well, their approach while being

more general, is also more complicated and involved. [20] also consider a special case of Problem 1, with \mathcal{C} being the family of cuts (i.e., the cooperative cut problem). Interestingly, they suggest an algorithm that is identical to PLA when f is a (low rank) sum of truncations (i.e., $\psi_i(x) = \min(x, a)$). For general sums of low-rank concave functions, they resort to the algorithms of [23, 8]. We provide a generic algorithm, which not only works for a much large class of constraints and functions, but also extends to the Problems 2, 3 and 4. Moreover, it is easy to see that our algorithms would also work for the more general problem of minimizing low rank sums of concave functions, over a solvable polytope.

Note that the complexity of PLA is polynomial in $\frac{1}{\epsilon}$, but exponential in k . Hence this makes sense only if k is a constant or is $O(\log n)$. If k is a constant (with respect to n), PLA is a fully polynomial time approximation scheme (FPTAS) [30]. If $k = O(\log n)$, then PLA is a polynomial-time approximation scheme (PTAS). This assumption is reasonable for many of the applications of Problems 1-4 (see details of this in the experiments section). Moreover, there are a number of ways one can speed up PLA. A very simple observation is that PLA is amenable to a distributive implementation via Map-Reduce. In particular, let $N = O(\frac{1}{\epsilon^k})$ denote the total number of computations of PLA (i.e., this is the number of times one performs an instance of Problems 1-4 with a modular function). All these can be performed in parallel on m processing systems. We output the best from each system to a central processor, which finds the optimal amongst these. The complexity of this distributive procedure is $O(NT/m + m)$, (where T is the complexity of using a modular function in the place of f in Problems 1-4), which improves the overall complexity by a factor of m .

In addition, we can also provide early stopping criterion and heuristics for speeding up PLA. One strategy of implementing PLA, is to start with $j_i = 1, \forall i = 1, 2, \dots, k$, and incrementally increase j_i in a coordinate ascent fashion. The following lemma gives a sufficient condition for stopping PLA.

Lemma 5. *Let $J = [j_1, j_2, \dots, j_k]$ be such that the corresponding solution \hat{X}_J satisfies $w_i(X_J) \in [b_i^{j_i}, b_i^{j_i+1}], \forall i$. Then \hat{X}_J is the (near) optimal solution for Problems 1, 2 and 3.*

The values of $w_i(X_J)$ also suggest the direction of the co-ordinate wise algorithm. For example, if $w_i(X_J) < b_i^{j_i}$, it suggests that the value of j_i be decreased. Similarly, if $w_i(X_J) > b_i^{j_i+1}$, its a sign that j_i be decreased. In this manner, one can define a greedy like heuristic to implement PLA [20], which picks for every coordinate, the slope which increases the objective value the most. Many of these heuristics have been considered in [20] in the case of cuts, and when the function class is low rank

sums of truncations. These heuristics are all polynomial in k , but are not guaranteed to obtain the optimal solutions. Moreover, in certain cases (for example, the case of cuts), one can do parametric versions, thereby solving a set of related problems simultaneously [20, 5].

4 Experiments

We next experimentally evaluate the performance of our methods. The utility of the constrained minimization algorithms for cooperative cuts have been investigated in [20]. In this paper, we consider the applications of cooperative image matching and sensor placement.

4.1 Cooperative Image matching

The problem of matching key-points in images, also called image correspondence, is an important problem in computer vision [26]. The simplest model for this problem constructs a matching with linear scores, i.e., a max bipartite matching [18], called a *linear assignment*. This model does not allow a representation of interaction between the pixels. For example, we see many obviously spurious matches in figure 4b. Many models try to capture this, via, for example via quadratic assignments [3]. Instead of just looking at the best linear assignment, the quadratic models try to incorporate pairwise constraints. This is also called graph matching.

We describe a new and different model here. First, we cluster key-points, separately in each of the two images, into k clusters. Figure 4a shows a particular clustering of an image into $k = 3$ groups. The clustering can be performed based on the pixel color map, or simply the distance of the key-points. That is, each image has k clusters. Let $\{V_i^{(1)}\}_{i=1}^k$ and $\{V_i^{(2)}\}_{i=1}^k$ be the two sets of clusters. We then compute the linear assignment problem, letting $\mathcal{M} \subseteq \mathcal{E}$ be the resulting maximum matching. We then partition the edge set $\mathcal{E} = \mathcal{E}_1 \cup \mathcal{E}_2 \cup \dots \mathcal{E}_k \cup \mathcal{E}'$ where $\mathcal{E}_i = \mathcal{M} \cap (V_\ell^{(1)} \times V_s^{(2)})$ for $\ell, s \in \{1, 2, \dots, k\}$ corresponding to the i 'th largest intersection, and $\mathcal{E}' = \{\mathcal{E} \setminus \bigcup_{i=1}^k \mathcal{E}_i\}$ are the remaining edges either that were not matched or that did not lie within a frequently associated pair of image key-point clusters. We then define a submodular function as follows:

$$f(S) = \sum_{i=1}^k \psi_i(w(S \cap \mathcal{E}_i)) + w(S \cap \mathcal{E}'), \quad (7)$$

which provides an additional discount to the edges $\{\mathcal{E}_i\}_{i=1}^k$ corresponding to key-points that were frequently associated in the initial pass. The problem of co-operative matching then becomes an instance of Problem 1 with the submodular function (over the edges) defined above, and a constraint that the edges

form a matching. Figures 4b and 4c shows how the submodular matchings improve over the simple bipartite matching, with $k = 3$. The minimum matching approach obtains many spurious matches between clusters (shown in red), while the cooperation described above reduces these spurious matches. The cooperative matching improves the performance over the modular method on these images by about 20%.

We also test the performance of our algorithms on the CMU House and Hotel dataset [3]. The house dataset has 111 images, while the hotel dataset has 101 images. We consider all possible pairs of images, with differences between the two images ranging from 0 : 10 : 90 in both cases. We consider three algorithms: PLA, SGA (both using Equation (7)) and the simple modular bipartite matching as a baseline (Mod). Again, we set $k = 3$. The results are shown in Figure 4(d-e) where we observe that PLA and SGA beat Mod by about 3 – 5% on average. Moreover, we also see that PLA, in general, outperforms SGA, thus showing how superior theoretical guarantees translate into better empirical performance. In PLA, we chose ϵ such that each concave function ψ_i has four break points. We observed, moreover, that setting lower values of ϵ does not improve the objective value in this application. We observe, moreover, that PLA also beats SGA in terms of objective value. In terms of running time, PLA and SGA are comparable. We do not compare the ellipsoidal approximation algorithm (EA) [7], mainly because it is too slow to run on real world problems. Moreover, this algorithm has been observed to perform comparably to the much simpler SGA [15]. While we considered the simple linear assignment as a baseline for the cooperative matching, it seems possible to embed this cooperation on more involved graph matching models as well.

4.2 Sensor Placement

We next consider an application of sensor placement. A number of natural models for this problem are forms of submodular maximization [21, 22]. A natural model, that performs very well in practice, is to maximize the mutual information $I(X_A; X_{V \setminus A})$, where A refers to the set of sensors chosen. [21, 22] investigate this in the setting of additive costs on the sensors. Often however, the costs are not additive in practice. In fact, very often, they are also submodular [22], and a natural model is,

$$f(X) = \sum_{i=1}^k \psi_i(c(X \cap S_i)) \quad (8)$$

where ψ_i s are concave, $c(j)$ is the cost of sensor j and S_1, \dots, S_k are groups of similar sensors. This was posed as an open problem in [22]. We can naturally pose this as instances of Problem 3, where

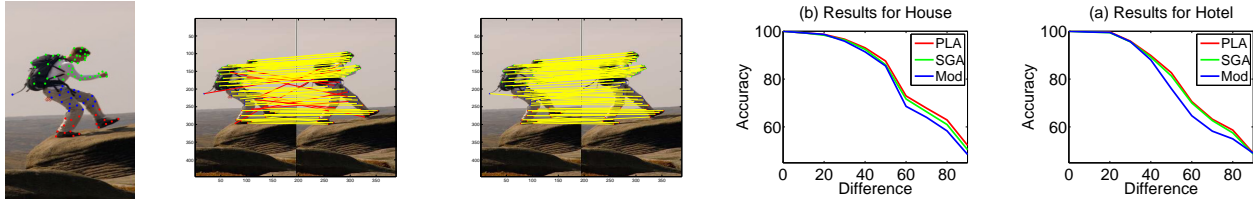


Figure 4: (a) (far left) shows the clustering used in cooperative matching, (b) and (c) show the results with bipartite matching and cooperative matching respectively, and (d) and (e) (far right) give the results on the [House](#) and [Hotel](#) dataset, showing PLA (this paper) achieves slightly better than SGA, and both submodular methods performing better than standard matching (Mod).

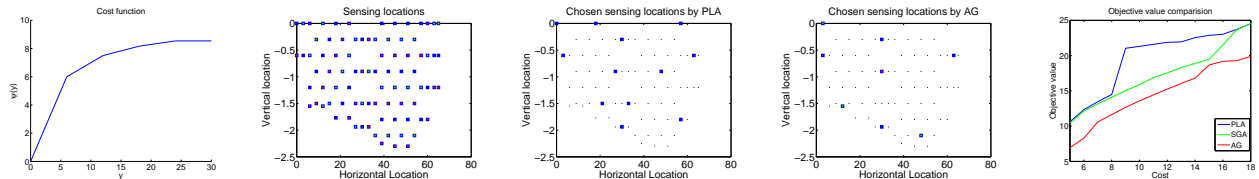


Figure 5: (a) Cost function $\psi_i(y)$ used, (b) set of locations with the three colors (blue, red, green) referring to the three types of sensors, (c) Sensors chosen by PLA, (d) Sensors chosen by AG and (e) plots of the objective value of different algorithms (please zoom in for details).

$g(X) = I(X_A; X_{V \setminus A})$ and $f(X)$ is the cost function above. Note that we could equivalently also express this as an instance of Problem 2 with a constraint on g while minimizing f .

We consider real world data of placing sensors to predict the pH values from the lake of Merced [21]. We also assume that the function f is piece-wise linear, shown in Figure 5a (far left). Figure 5b shows the locations (horizontal and vertical). We assume that there are three kinds of locations, shown in blue, green and red colors respectively, and the costs of placing sensors in the same kind of location is discounted. Correspondingly, we assume the cost function is an instance of function Equation (8) with $k = 3$. For simplicity, we assume also that all three types of sensor locations have the same coverage model (though, in general, it would make sense for them to have different models for coverage, based on their type). Under this assumption, the optimal configuration would tend to be spatially diverse, yet cooperative (in the sense, that the same type of sensors would be chosen).

We compare three algorithms: PLA, and SGA (both on Problem 3), and a simple cost agnostic greedy algorithm (AG), which ignores the cost function f , and greedily adds sensors. Figure 5c shows the sensors chosen by PLA (the cost sensitive one), and Figure 5d shows the choices of AG (the cost agnostic one). While both have the same cost budget, the cost agnostic one does not utilize the discounts of placing sensors in similar locations, and correspondingly, places fewer sensors. The cost sensitive algorithms (PLA and SGA) on the other hand, simultaneously achieve coverage, while making use of the discounts. Figure 5e plots the

objective functions attained by the three algorithms. We see that both PLA and SGA, outperform the agnostic greedy algorithm. Moreover, PLA also performs better than SGA. Note that the function f , used in this case is piece-wise linear, and correspondingly PLA is exact in this case.

5 Conclusions

In this paper, we investigated a new class of algorithms for various forms of constrained submodular programs, with a special subclass of submodular cost functions. We focus on problems that for the general class of submodular functions are hard, and yet occur naturally in many applications. We showed that when we restrict the class of functions to low rank sums of concave over modular functions, we can obtain significantly improved worst case theoretical results. We also complemented our results with experimental results in sensor placement and image correspondence. An immediate open question is whether there are similar algorithms for other rich and useful subclasses of submodular functions. In particular, it would be interesting if one can remove the low rank assumption, and provide tighter approximation algorithms for general sums of concave over modular functions, which would be very powerful.

This material is based upon work supported by the National Science Foundation under Grant No. (IIS-1162606), as well as a Google and a Microsoft award. This work was also 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] W. Bai, R. Iyer, K. Wei, and J. Bilmes. Algorithms for optimizing the ratio of submodular functions. In *International Conference on Machine Learning*, pages 2751–2759, 2016.
- [2] Y. Boykov and V. Kolmogorov. An experimental comparison of min-cut/max-flow algorithms for energy minimization in vision. *TPAMI*, 26(9):1124–1137, 2004.
- [3] T. S. Caetano, J. J. McAuley, L. Cheng, Q. V. Le, and A. J. Smola. Learning graph matching. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 31(6):1048–1058, 2009.
- [4] S. Fujishige. *Submodular functions and optimization*, volume 58. Elsevier Science, 2005.
- [5] S. Fujishige and S. Iwata. Minimizing a submodular function arising from a concave function. *Discrete applied mathematics*, 92(2):211–215, 1999.
- [6] G. Goel, P. Tripathi, and L. Wang. Combinatorial problems with discounted price functions in multi-agent systems. In *FSTTCS*, 2010.
- [7] M. Goemans, N. Harvey, S. Iwata, and V. Mirrokni. Approximating submodular functions everywhere. In *SODA*, pages 535–544, 2009.
- [8] V. Goyal and R. Ravi. An fptas for minimizing a class of low-rank quasi-concave functions over a convex set. *Operations Research Letters*, 41(2):191–196, 2013.
- [9] S. Iwata and K. Nagano. Submodular function minimization under covering constraints. In *In FOCS*, pages 671–680. IEEE, 2009.
- [10] R. Iyer and J. Bilmes. Algorithms for approximate minimization of the difference between submodular functions, with applications. In *UAI*, 2012.
- [11] R. Iyer and J. Bilmes. The submodular Bregman and Lovász-Bregman divergences with applications. In *NIPS*, 2012.
- [12] R. Iyer and J. Bilmes. Submodular Optimization with Submodular Cover and Submodular Knapsack Constraints. In *NIPS*, 2013.
- [13] R. Iyer, S. Jegelka, and J. Bilmes. Curvature and Optimal Algorithms for Learning and Minimizing Submodular Functions. In *Neural Information Processing Society (NIPS)*, 2013.
- [14] R. Iyer, S. Jegelka, and J. Bilmes. Curvature and Optimal Algorithms for Learning and Optimization of Submodular Functions: Extended arxiv version, 2013.
- [15] R. Iyer, S. Jegelka, and J. Bilmes. Fast Semidifferential based Submodular function optimization. In *ICML*, 2013.
- [16] S. Jegelka and J. A. Bilmes. Approximation bounds for inference using cooperative cuts. In *ICML*, 2011.
- [17] S. Jegelka and J. A. Bilmes. Submodularity beyond submodular energies: coupling edges in graph cuts. In *CVPR*, 2011.
- [18] S. Jegelka, A. Kapoor, and E. Horvitz. An interactive approach to solving correspondence problems. *International Journal of Computer Vision*, pages 1–10, 2013.
- [19] J. A. Kelner and E. Nikolova. On the hardness and smoothed complexity of quasi-concave minimization. In *Foundations of Computer Science, 2007. FOCS'07. 48th Annual IEEE Symposium on*, pages 472–482. IEEE, 2007.
- [20] P. Kohli, A. Osokin, and S. Jegelka. A principled deep random field for image segmentation. In *CVPR*, 2013.
- [21] A. Krause and C. Guestrin. Optimizing sensing: From water to the web. Technical report, DTIC Document, 2009.
- [22] A. Krause, A. Singh, and C. Guestrin. Near-optimal sensor placements in Gaussian processes: Theory, efficient algorithms and empirical studies. *JMLR*, 9:235–284, 2008.
- [23] S. Mittal and A. S. Schulz. An fptas for optimizing a class of low-rank functions over a polytope. *Mathematical Programming*, 141(1-2):103–120, 2013.
- [24] G. Nemhauser and L. Wolsey. Best algorithms for approximating the maximum of a submodular set function. *Mathematics of Operations Research*, 3(3):177–188, 1978.
- [25] E. Nikolova. Approximation algorithms for offline risk-averse combinatorial optimization, 2010.
- [26] A. S. Ogale and Y. Aloimonos. Shape and the stereo correspondence problem. *International Journal of Computer Vision*, 65(3):147–162, 2005.
- [27] C. Qian, J.-C. Shi, Y. Yu, K. Tang, and Z.-H. Zhou. Optimizing ratio of monotone set functions. In *IJCAI*, 2017.
- [28] M. Sviridenko. A note on maximizing a submodular set function subject to a knapsack constraint. *Operations Research Letters*, 32(1):41–43, 2004.
- [29] Z. Svitkina and L. Fleischer. Submodular approximation: Sampling-based algorithms and lower bounds. In *FOCS*, pages 697–706, 2008.

- [30] V. V. Vazirani. *Approximation algorithms*.
springer, 2004.
- [31] L. A. Wolsey. An analysis of the greedy algorithm for the submodular set covering problem.
Combinatorica, 2(4):385–393, 1982.