

On the Benefits of Progressively Increasing Sampling Sizes in Stochastic Greedy Weak Submodular Maximization

Abolfazl Hashemi ^{1b}, Member, IEEE, Haris Vikalo ^{1b}, Senior Member, IEEE, and Gustavo de Veciana ^{1b}, Fellow, IEEE

Abstract—Many problems in signal processing and machine learning can be formalized as weak submodular optimization tasks. For such problems, a simple greedy algorithm (GREEDY) is guaranteed to find a solution achieving the objective with a value no worse than $1 - e^{-1/c}$ of the optimal, where c is the multiplicative weak-submodularity constant. Due to the high cost of querying large-scale systems, the complexity of GREEDY becomes prohibitive in contemporary applications. In this work, we study the tradeoff between performance and complexity when one resorts to random sampling strategies to reduce the query complexity of GREEDY. Specifically, we quantify the effect of uniform sampling strategies on GREEDY's performance through two metrics: (i) asymptotic probability of identifying an optimal subset, and (ii) suboptimality with respect to the optimal solution. The latter implies that uniform sampling strategies with a fixed sampling size achieve a non-trivial approximation factor; however, we show that with overwhelming probability, these methods fail to find the optimal subset. Our analysis shows that the failure of uniform sampling strategies with fixed sample size can be circumvented by successively increasing the size of the search space. Building upon this insight, we propose a simple progressive stochastic greedy algorithm and study its approximation guarantees. Moreover, we demonstrate effectiveness of the proposed method in dimensionality reduction applications and feature selection tasks for clustering and object tracking.

Index Terms—Weak submodular optimization, greedy algorithms, randomized algorithms, subset selection.

I. INTRODUCTION

WE STUDY the problem of maximizing non-decreasing weak submodular functions under a cardinality constraint in large-scale settings. The well-known GREEDY algorithm [2] selects the solution set by sequentially identifying

elements with the largest marginal contribution; the algorithm achieves a $1 - 1/e$ worst-case approximation for monotone submodular functions [2], the tightest guarantee for any algorithm that can evaluate the objective function on only polynomially many inputs. Although GREEDY achieves the optimal approximation factor, the computational cost of doing so is expensive for large-scale problems. This motivates the search for approximation schemes capable of accelerating the optimization without significant sacrifice of accuracy.

Recently, greedy algorithms that utilize random sampling while restricting size of the searched set to a fixed value have been proposed; for a problem with cardinality constraint k and the ground set of size m , such methods incur complexity of only $\mathcal{O}(m \log \frac{1}{\epsilon})$ [3]–[6]. In expectation, the fixed search size methods achieve a constant factor approximation of $1 - 1/e - \epsilon$, nearly matching the worst-case performance guarantee of GREEDY while providing a computational gain of $\mathcal{O}(\frac{k}{\log \frac{1}{\epsilon}})$. Motivated by the success of greedy search schemes in practical settings, in this paper we investigate the impact of the search space size on their performance and study performance-complexity tradeoffs in maximizing a (weak) submodular function. To this end, we consider two criteria: (1) the ability of greedy algorithms with uniform sampling to exactly identify the optimal solution to a (weak) submodular maximization problem, and (2) the degree of suboptimality of the selected solution with respect to the optimal value.

Our first contribution, formalized in Theorem 2, states that as the size of the ground set and cardinality constraint increase, randomized greedy schemes with a restricted search space with overwhelming probability fail to successfully identify the optimal subset. This in turn implies that, while there may be scenarios where GREEDY can identify the optimal subset with high-probability (e.g., the task of sparse recovery), there is an unbounded gap between the exact identification capability of GREEDY and randomized schemes with a restricted search space.

Aiming to overcome the above limitation, as part of our next contribution we establish that having an increasing schedule of sampling set sizes is unavoidable. Building on this insight, we propose a new algorithm that we refer to as *Progressive Stochastic Greedy* (PSG) and analyze its achievable performance. In particular, we show that PSG attains improved worst-case approximation factor, both on expectation and with high probability, compared to randomized greedy schemes with a fixed sampling size.

Manuscript received 20 October 2021; revised 26 May 2022 and 21 July 2022; accepted 22 July 2022. Date of publication 29 July 2022; date of current version 15 August 2022. The associate editor coordinating the review of this manuscript and approving it for publication was Dr. Ketan Rajawat. This work was supported by NSF under Grant ECCS-1809327. This work was presented in part at the IEEE International Conference on Acoustics, Speech and Signal Processing, Toronto, Ontario, Canada, June 2021 [DOI: 10.1109/ICASSP39728.2021]. (Corresponding author: Abolfazl Hashemi.)

Abolfazl Hashemi is with the School of Electrical and Computer Engineering, Purdue University, West Lafayette, IN 47907 USA (e-mail: abolfazl@purdue.edu).

Haris Vikalo and Gustavo de Veciana are with the Department of Electrical and Computer Engineering, University of Texas at Austin, Austin, TX 78712 USA (e-mail: hvikalo@ece.utexas.edu; deveciana@utexas.edu).

This article has supplementary downloadable material available at <https://doi.org/10.1109/TSP.2022.3195089>, provided by the authors.

Digital Object Identifier 10.1109/TSP.2022.3195089

Finally, we consider the application of the proposed scheme to two subset selection problems, namely, column subset selection for subspace clustering and observation selection for extended object tracking in automated driving. Our results demonstrate the efficacy of the proposed scheme in reducing the computational costs of subset selection with negligible performance drop.

A. Related Work

Submodularity is a property of set functions with desirable theoretical and practical implications relevant to many problems in combinatorial optimization. For instance, submodular maximization applies to many well-known problems such as facility location, coverage problems, and maximum weighted matching in discrete optimization [7] as well as active learning, influence maximization, and information gathering in machine learning [8]–[10]. In such problems, the goal is to maximize a monotonically increasing submodular function subject to a linear matroid, or a cardinality constraint.

The objective function in some applications, e.g., sparse support selection and observation selection [4], [11]–[13], is not necessarily a submodular function; rather, one deals with *weakly* submodular objectives that resemble diminishing return property of submodular functions.

Recent advances in information systems have brought forth unprecedented amounts of data in many settings, including contemporary weak submodular maximization problems. Given a cardinality constraint k and a ground set of size m , the classical GREEDY algorithm for monotone weak submodular maximization that enjoys an optimal $1 - 1/e$ constant factor approximation [2] requires $\mathcal{O}(mk)$ function evaluations. Therefore, in data intensive applications where function evaluation is expensive, running GREEDY may be infeasible. To this end, there have been recent efforts to exploit strong theoretical guarantees of GREEDY while improving on its complexity via resorting to either distributed and parallel computing schemes, or methods for reducing the cost-per-iteration of GREEDY. Among the former, there is a growing line of work to design algorithms with sublinear *adaptivity* [14]–[17]. The concept of adaptivity is heavily studied in computer science and optimization; informally, adaptivity characterizes efficiency of parallel computation of an algorithm. The focus of this paper, however, is on the latter, i.e., centralized schemes – distributed weak submodular maximization methods are complementary to our study. Nevertheless, our analysis and the proposed algorithm may be deployed to aid distributed methods that rely on GREEDY, potentially extending their utility.

The LAZY-GREEDY algorithm [18] exploits the notion of submodularity to decrease the number of function evaluations of each iteration of GREEDY without sacrificing its performance. However, similar to GREEDY, LAZY-GREEDY incurs $\mathcal{O}(mk)$ function evaluations. Moreover, it cannot be employed in weak submodular maximization problems. More recently, Badaniyuru and Vondrak [19] proposed a randomized scheme that achieves a worst case approximation factor of $1 - 1/e - \epsilon$ while using $\mathcal{O}(\frac{m}{\epsilon} \log \frac{m}{\epsilon})$ evaluations. Motivated by this work, Mirzasoleyman et al. [3] proposed STOCHASTIC-GREEDY that achieves a worst case approximation factor of $1 - 1/e - \epsilon$ while using

$\mathcal{O}(m \log \frac{1}{\epsilon})$ function evaluations. Further discussion on greedy algorithms with random sampling is deferred to Section II-D. To our knowledge, no prior works study the effect of the sampling size on the success probability of the greedy algorithms with random sampling.

Sparse reconstruction and sparse support selection tasks belong to a class of cardinality-constrained weak submodular maximization problems where the exact identification of the optimal subset is of critical importance. In sparse support selection, the goal is to identify the support of a high dimensional vector (e.g., an image or a signal), i.e., the collection of nonzero components of the data, from a relatively small number of measurements. In such settings, GREEDY satisfies a general constant factor approximation guarantee as shown by [2], [13]. However, by exploiting the underlying structural properties of the measurement model in sparse support selection, one can establish conditions under which GREEDY *exactly* identifies the optimal subset. To this end, necessary and sufficient conditions for exact identification via GREEDY have been established by relying on various analysis techniques including those based on restricted isometry [20]–[22] and mutual incoherence [23]–[25] properties. In particular, when the measurements are randomly generated, Tropp and Gilbert [11] show that GREEDY enjoys an optimal sample complexity bound outlined by [26]. In the supplementary, we explore an application of the proposed algorithm to this task.

A related task of column subset selection (CSS) has received considerable attention in recent years due to its broad applicability, interpretability, and provably-guaranteed performance [27]–[29]. CSS is a constrained low-rank-approximation problem that seeks to approximate a data matrix (e.g., a matrix having data points in its rows and features in its columns) by projecting it onto a space spanned by only a few of its columns. While similar to the general low-rank approximation problem, CSS possesses certain distinguishing characteristics. First, since CSS is an unsupervised method and does not require labeled data, it can be applied efficiently to the scenarios where labeled data is sparse while unlabeled data is abundant. Second, to learn interpretable models in applications where the decision is made via a data-driven algorithm (e.g., hiring and education), it is of critical importance to keep the semantic interpretation of the features intact. This can be ensured by selecting a subset of available features as opposed to generating new features via an arbitrary function of the input features. Finally, compared to PCA or other methods that require a matrix-matrix multiplication to project input features into a reduced space during inference time, solution of CSS feature selection problem can be applied efficiently during inference as CSS only requires selecting a subset of feature values from a new instance vector.

An efficient approach to CSS relies on a greedy algorithm that selects the most representative subset of columns in an iterative fashion by greedily optimizing the reconstruction error [30], [31]. The iterative procedure of the greedy scheme is readily implemented in practice, and its often strong performance is complemented with rigorous theoretical guarantees [32], [33]. However, running the greedy scheme can be computationally expensive for large datasets. This is because if the goal is to

select k out of the m available columns, in each of k iterations of the greedy scheme one needs to evaluate marginal contribution of $\mathcal{O}(n)$ columns. Although computational costs can be reduced using the so-called lazy evaluations [34], the worst case number of function evaluations of the greedy scheme is $\mathcal{O}(nk)$. In our work, we apply the proposed PSG algorithm to improve the computational efficiency of greedy CSS schemes.

B. Organization

The rest of the paper is organized as follows. In Section II, we introduce the notation and review relevant concepts from weak submodular maximization. In Section III, we study the conditions for successful identifications of the optimal subset via uniform sampling strategies. In particular, we present our first contribution which establishes the failure of methods with a fixed sampling size to exactly identify the optimal subset, and argue the necessity of having an increasing schedule of sampling sizes; this insight leads to the proposed PSG algorithm. Section IV presents the analysis of PSG and provides bounds on its suboptimality performance. In Section V, we consider applications of PSG in subset selection problems. Finally, concluding remarks are provided in Section VI.

II. BACKGROUND

In this section, we introduce the notation and provide an overview of relevant concepts from submodular optimization as well as the GREEDY algorithm and methods with uniform sampling.

A. Notation

Italic letters represent scalars and numerical constants, e.g., α and C . We use calligraphic letters to denote sets, e.g., \mathcal{S} . Bold capital letters denote matrices, e.g., \mathbf{A} , while bold lowercase letters represent column vectors, e.g., \mathbf{a} . Finally, $\mathbb{I}(\cdot)$ denotes the indicator function of its argument.

B. Submodular Maximization

Definition 1: Set function $f : 2^{\mathcal{X}} \rightarrow \mathbb{R}$ is monotone if $f(\mathcal{S}) \leq f(\mathcal{T})$ for all $\mathcal{S} \subseteq \mathcal{T} \subseteq \mathcal{X}$.

Definition 2: Set function $f : 2^{\mathcal{X}} \rightarrow \mathbb{R}$ is submodular if

$$f(\mathcal{S} \cup \{j\}) - f(\mathcal{S}) \geq f(\mathcal{T} \cup \{j\}) - f(\mathcal{T}) \quad (1)$$

for all subsets $\mathcal{S} \subseteq \mathcal{T} \subseteq \mathcal{X}$ and $j \in \mathcal{X} \setminus \mathcal{T}$. The term $f_j(\mathcal{S}) = f(\mathcal{S} \cup \{j\}) - f(\mathcal{S})$ is the marginal value of adding element j to set \mathcal{S} .

Given a monotone non-decreasing set function $f : 2^{\mathcal{X}} \rightarrow \mathbb{R}$ with $f(\emptyset) = 0$, we are interested in solving the combinatorial optimization problem

$$\begin{aligned} & \underset{\mathcal{S}}{\text{maximize}} && f(\mathcal{S}) \\ & \text{subject to} && \mathcal{S} \subset \mathcal{X}, \quad |\mathcal{S}| \leq k, \end{aligned} \quad (2)$$

which we denote by $\mathcal{P}(m, k)$, where $|\mathcal{X}| = m$. By a reduction to the well-known set cover problem, the combinatorial optimization (2) can be shown to be NP-hard [7], [35]. It has been shown

in [2] that if $f(\cdot)$ is monotone and submodular, a simple greedy algorithm that iteratively selects an element with the highest marginal gain (see Algorithm 1) satisfies the optimal $1 - 1/e$ worst case approximation ratio.

C. Weak Submodularity

In many problems, the objective function is not submodular but under certain conditions it behaves similarly. Such functions are called *weakly* submodular and the extent of their proximity to submodularity is captured using the following parameters.

Definition 3: The multiplicative weak-submodularity constant of a monotone non-decreasing function f is defined as

$$c_f = \max_{(\mathcal{S}, \mathcal{T}, i) \in \tilde{\mathcal{X}}} f_i(\mathcal{T}) / f_i(\mathcal{S}), \quad (3)$$

where $\tilde{\mathcal{X}} = \{(\mathcal{S}, \mathcal{T}, i) | \mathcal{S} \subseteq \mathcal{T} \subseteq \mathcal{X}, i \in \mathcal{X} \setminus \mathcal{T}\}$.

The multiplicative weak-submodularity constant [36]–[38] is a closely related concept to submodularity and essentially quantifies how close the set function is to being submodular. It is worth noting that a set function $f(\mathcal{S})$ is submodular if and only if its multiplicative weak-submodularity constant satisfies $c_f \leq 1$ [4], [13], [39].

A similar notion of weak submodularity is the additive weak-submodularity constant defined below [36]–[38].

Definition 4: The additive weak-submodularity constant of a monotone non-decreasing function f is defined as

$$\epsilon_f = \max_{(\mathcal{S}, \mathcal{T}, i) \in \tilde{\mathcal{X}}} f_i(\mathcal{T}) - f_i(\mathcal{S}), \quad (4)$$

where $\tilde{\mathcal{X}} = \{(\mathcal{S}, \mathcal{T}, i) | \mathcal{S} \subseteq \mathcal{T} \subseteq \mathcal{X}, i \in \mathcal{X} \setminus \mathcal{T}\}$.

Note that when $f(\mathcal{S})$ is submodular, its additive weak-submodularity constant satisfies $\epsilon_f \leq 0$.

For a monotone function with bounded additive and multiplicative weak-submodularity constants (WSCs) we have the following proposition (see, e.g., [38]).

Proposition 1: Let c_f and e_f be the multiplicative and additive weak-submodularity constants of $f(\mathcal{S})$, a monotone non-decreasing function with $f(\emptyset) = 0$, respectively. Let \mathcal{S} and \mathcal{T} be any subsets such that $\mathcal{S} \subset \mathcal{T} \subseteq \mathcal{X}$ with $|\mathcal{T} \setminus \mathcal{S}| = r$. Then, it holds that

$$f(\mathcal{T}) - f(\mathcal{S}) \leq \frac{1}{r} (1 + (r-1)c_f) \sum_{j \in \mathcal{T} \setminus \mathcal{S}} f_j(\mathcal{S}), \quad (5)$$

and

$$f(\mathcal{T}) - f(\mathcal{S}) \leq (r-1)e_f + \sum_{j \in \mathcal{T} \setminus \mathcal{S}} f_j(\mathcal{S}). \quad (6)$$

It is worth pointing out other weak submodularity notions such as those in [4], [13], [36], [39] that depending on the application may simplify the derivation of the approximation bounds (see e.g., [37], [38], [40]–[42]).

Using notion of weak-submodularity, one can extend the theoretical results of [2] for GREEDY to the case of weak submodular functions [13]. In particular, the following proposition holds (see, e.g. [38]).

Algorithm 1: GREEDY.

-
- 1: **Input:** Weak submodular function f , ground set \mathcal{X} , number of elements to be selected k .
 - 2: **Output:** Subset $\mathcal{S}_g \subseteq \mathcal{X}$ with $|\mathcal{S}_g| = k$.
 - 3: Initialize $\mathcal{S}_g^{(0)} = \emptyset$
 - 4: **for** $i = 0, \dots, k - 1$ **do**
 - 5: $j_s \in \operatorname{argmax}_{j \in \mathcal{X}} f_j(\mathcal{S}_g^{(i)})$
 - 6: $\mathcal{S}_g^{(i+1)} = \mathcal{S}_g^{(i)} \cup \{j_s\}$
 - 7: **end for**
 - 8: **return** $\mathcal{S}_g = \mathcal{S}_g^{(k)}$.
-

Theorem 1: Let c_f and e_f be the multiplicative and additive weak-submodularity constants of $f(\mathcal{S})$, a monotone non-decreasing function with $f(\emptyset) = 0$. Let $\mathcal{S}_g \subseteq \mathcal{X}$ with $|\mathcal{S}_g| \leq k$ be the subset selected when maximizing $f(\mathcal{S})$ subject to a cardinality constraint via the greedy observation selection scheme, and let \mathcal{S}^* denote the optimal subset. Then

$$f(\mathcal{S}_g) \geq \left(1 - e^{-\frac{1}{c}}\right) f(\mathcal{S}^*), \quad (7)$$

where $c = \max\{c_f, 1\}^1$ and

$$f(\mathcal{S}_g) \geq \left(1 - \frac{1}{e}\right) (f(\mathcal{S}^*) - (k - 1)e_f). \quad (8)$$

The approximation results in Proposition 1 imply that if the objective function is monotone and weak submodular, the greedy selection scheme which in each iteration selects an element with the highest marginal gain finds a solution that is close to the optimal.

D. GREEDY with Random Sampling

If $|\mathcal{X}| = m$, in each of k iterations of GREEDY one needs to find the marginal gain of $\mathcal{O}(m)$ elements. This is a computationally intensive procedure when the involved datasets are large. Although the computational costs can be reduced using the so-called lazy evaluations [34], the worst case number of function evaluations of GREEDY is $\mathcal{O}(mk)$. The prohibitive complexity of GREEDY for large-scale datasets has motivated the design of more efficient schemes for weak submodular maximization. A body of work typically referred to as STOCHASTIC-GREEDY algorithms, aims to reduce the number of function evaluations by restricting the search domain in each iteration of the greedy selection procedure [3]–[6] (see Algorithm 2 for a general template). Specifically, instead of evaluating the marginal gain of $\mathcal{O}(m)$ elements, in the i th iteration of greedy-with-random-sampling one selects a subset $\mathcal{R}^{(i)} \subseteq \mathcal{X}$ by uniformly at random sampling r_i elements and only evaluates marginal gains of the elements in $\mathcal{R}^{(i)}$. References [3]–[6] explore the case where $\mathcal{R}^{(i)} = \mathcal{R}$ is fixed and $r_i = r \frac{m}{k} \log \frac{1}{\epsilon}$, where the parameter ϵ , $0 < e^{-k} \leq \epsilon \leq e^{-\frac{k}{m}} < 1$, determines the size of the search

¹Henceforth, we assume $c_f \geq 1$ to emphasize that the objective function is typically weak submodular.

Algorithm 2: GREEDY With Random Sampling.

-
- 1: **Input:** Weak submodular function f , ground set \mathcal{X} , number of elements to be selected k , search space schedule $\{r_i\}_{i=0}^{k-1}$.
 - 2: **Output:** Subset $\mathcal{S}^{(k)} \subseteq \mathcal{X}$ with $|\mathcal{S}^{(k)}| = k$.
 - 3: Initialize $\mathcal{S}^{(0)} = \emptyset$
 - 4: **for** $i = 0, \dots, k - 1$ **do**
 - 5: Form $\mathcal{R}^{(i)}$ by sampling $\min(r_i, m)$ elements from $\mathcal{X} \setminus \mathcal{S}^{(i)}$ uniformly at random.
 - 6: $j_s \in \operatorname{argmax}_{j \in \mathcal{R}^{(i)}} f_j(\mathcal{S}^{(i)})$
 - 7: $\mathcal{S}^{(i+1)} = \mathcal{S}^{(i)} \cup \{j_s\}$
 - 8: **end for**
-

domain $\mathcal{R}^{(i)}$ and thus controls the number of function evaluations in each iteration.² It turns out that with this choice, the complexity of greedy-with-random-sampling is $\mathcal{O}(m \log \frac{1}{\epsilon})$ and that it selects a subset \mathcal{S}_{sg} such that

$$\mathbb{E}[f(\mathcal{S}_{sg})] \geq (1 - 1/e - \epsilon) f(\mathcal{S}^*), \quad (9)$$

given that f in (2) is submodular [3]. This approximation ratio is derived under the simplifying assumption that the sequence of random subsets $\{\mathcal{R}_{sg}^{(i)}\}_{i=0}^{k-1}$ is constructed via sampling with replacement. Khanna et al. [41] analyze STOCHASTIC-GREEDY for weak submodular functions and show that it achieves an expected $1 - e^{-1/c} - \epsilon$ worst case approximation ratio. For the setting where $\{\mathcal{R}_{sg}^{(i)}\}_{i=0}^{k-1}$ are constructed via sampling without replacement, this approximation ratio is improved to $1 - e^{-1/c} - \epsilon^\beta/c$, where $\beta = 1 + \mathcal{O}(1/k)$ [5]. Under a specific set of assumptions on the marginal gain of elements selected in each iteration, [6] have shown that a similar result holds not only on expectation, but with high probability.

III. ANALYSIS OF SUCCESS PROBABILITY

We start the analysis by studying the probability of successfully identifying the optimal subset \mathcal{S}^* . This notion is formalized in the following definition.

Definition 5: Let ALG be an approximation algorithm for the weak submodular optimization problem (2) with a unique solution \mathcal{S}^* . Let \mathcal{S}_{alg} be the output of ALG. Then, ALG successfully identifies \mathcal{S}^* if $\mathcal{S}_{alg} = \mathcal{S}^*$. Furthermore, the probability of success of ALG is defined as $\Pr(\mathcal{S}_{alg} = \mathcal{S}^*)$.

Our goal in this section is to quantify the impact of r_i on the performance of Algorithm 2 by analyzing its probability of success. To successfully identify \mathcal{S}^* , in each iteration of Algorithm 2 at least one new (not previously selected) element of \mathcal{S}^* should be present in the randomly selected subset $\mathcal{R}^{(i)}$. More formally, if $\mathcal{S}^{(i)}$ denotes the subset of elements selected by Algorithm 2 before executing the i th iteration, $i = 0, \dots, k - 1$, the set $\mathcal{R}^{(i)} \cap (\mathcal{S}^* \setminus \mathcal{S}^{(i)})$ should be nonempty. This, however, is not sufficient – to find the optimal solution, Algorithm 2 must in each iteration select elements from $\mathcal{R}^{(i)} \cap (\mathcal{S}^* \setminus \mathcal{S}^{(i)})$. Since $|\mathcal{S}^*| = k$ and since in each iteration Algorithm 2 selects

²In this paper, for simplicity of presentation we assume $\log \frac{1}{\epsilon}$ and r are integer-valued quantities.

one element, if there exists an $i \in [k]$ such that $\mathcal{S}^{(i)} \not\subseteq \mathcal{S}^*$, then Algorithm 2 fails to identify \mathcal{S}^* . This informal argument underlies Lemma 1 below which will allow us to characterize the success probability.

Lemma 1: Suppose the optimal solution of (2) is unique. Let $\mathcal{S}^{(k)}$ denote the subset selected by Algorithm 2, and let $\mathcal{R}^{(i)}$ denote the randomly selected search space of Algorithm 2 in the i th iteration. Then, it holds that

$$\Pr\left(\mathcal{S}^{(k)} = \mathcal{S}^*\right) = \prod_{i=0}^{k-1} p^{(i)} \prod_{i=0}^{k-1} q^{(i)}, \quad (10)$$

where

$$p^{(i)} = \Pr\left(\mathcal{R}^{(i)} \cap (\mathcal{S}^* \setminus \mathcal{S}^{(i)}) \neq \emptyset \mid \mathcal{S}^{(i)} \subset \mathcal{S}^*, |\mathcal{S}^{(i)}| = i\right), \quad (11)$$

and

$$q^{(i)} = \Pr\left(\mathcal{S}^{(i+1)} \subset \mathcal{S}^* \mid \mathcal{R}^{(i)} \cap (\mathcal{S}^* \setminus \mathcal{S}^{(i)}) \neq \emptyset, |\mathcal{S}^{(i)}| = i\right). \quad (12)$$

Proof: Let $\mathcal{A}^{(i)}$ denote the event $\{\mathcal{S}^{(i+1)} \cap \mathcal{S}^{(i)} \neq \mathcal{S}^{(i)}, \mathcal{S}^{(i+1)} \subseteq \mathcal{S}^*\}$. Then the probability of success of Algorithm 2 can be expressed as

$$\begin{aligned} \Pr\left(\mathcal{S}^{(k)} = \mathcal{S}^*\right) &= \Pr\left(\mathcal{A}^{(k)}\right) \\ &= \prod_{i=0}^{k-1} \Pr\left(\mathcal{A}^{(i)} \mid \mathcal{A}^{(i-1)}\right). \end{aligned} \quad (13)$$

Note that $\mathcal{A}^{(i)}$ can equivalently be written as

$$\mathcal{A}^{(i)} = \{\mathcal{R}^{(i)} \cap (\mathcal{S}^* \setminus \mathcal{S}^{(i)}) \neq \emptyset, \mathcal{S}^{(i+1)} \subseteq \mathcal{S}^*\}. \quad (14)$$

To see this, recall $\mathcal{A}^{(i)}$ denotes the event $\{\mathcal{S}^{(i+1)} \cap \mathcal{S}^{(i)} \neq \mathcal{S}^{(i)}, \mathcal{S}^{(i+1)} \subseteq \mathcal{S}^*\}$. For this to happen, the sampling set at iteration i should have at least one element from \mathcal{S}^* that is not in $\mathcal{S}^{(i)}$. Conversely, if $\mathcal{S}^{(i+1)} \cap \mathcal{S}^{(i)} \neq \mathcal{S}^{(i)}$ and $\mathcal{S}^{(i+1)} \subseteq \mathcal{S}^*$, then, there must have been an element from \mathcal{S}^* in the sampling set which is new. By further conditioning, this can be re-written as

$$\Pr\left(\mathcal{S}^{(k)} = \mathcal{S}^*\right) = \prod_{i=0}^{k-1} p^{(i)} \prod_{i=0}^{k-1} q^{(i)}, \quad (15)$$

where $p^{(i)}$ and $q^{(i)}$ are given by (11) and (12), respectively. ■

Lemma 1 demonstrates that the probability of success of Algorithm 2 is product of two terms: (i) $p = \prod_{i=0}^{k-1} p^{(i)}$ that characterizes the likelihood of the event \mathcal{E}_1 that $\mathcal{R}^{(i)}$ contains at least one new element of \mathcal{S}^* for all $i = 0, \dots, k-1$, (ii) $q = \prod_{i=0}^{k-1} q^{(i)}$ that characterizes the likelihood of the event \mathcal{E}_2 of selecting one of the elements in the nonempty intersection of search space $\mathcal{R}^{(i)}$ and $\mathcal{S}^* \setminus \mathcal{S}^{(i)}$. The events \mathcal{E}_1 and \mathcal{E}_2 collectively are necessary and sufficient conditions for exact identification of the optimal subset.

The first probability, $p = \prod_{i=0}^{k-1} p^{(i)}$, is of particular interest as it can be thought of as being a general upper bound on the probability of success. Following this idea, next we establish an upper bound on the asymptotic probability of success for large-scale problems and show that some variants with overwhelming probability fail to successfully identify the optimal subset.

Theorem 2: Consider a sequence of optimization problems $\mathcal{P}(m, k)$ in (2) with increasingly higher dimensions, i.e., the setting where $m, k \rightarrow \infty$, $m > k$. Let ALG denote a variant of Algorithm 2 with a restricted uniform search space $\mathcal{R}^{(i)} \subset [m]$ having cardinality r_i such that $\max_i r_i < m - k$ and $\limsup_{m, k \rightarrow \infty} r_i/m = 0$, for all i . Then the probability that ALG succeeds on $\mathcal{P}(m, k)$ goes to zero, i.e.,³

$$\limsup_{m, k \rightarrow \infty} \Pr\left(\mathcal{S}_{alg}^{(k)} = \mathcal{S}^*\right) = 0. \quad (16)$$

Proof: First, assume ALG uses sampling with replacement to construct $\mathcal{R}_{alg}^{(i)}$. We can compute $p_{alg}^{(i)}$ according to

$$\begin{aligned} p_{alg}^{(i)} &= 1 - \Pr\left(\mathcal{R}_{alg}^{(i)} \cap (\mathcal{S}^* \setminus \mathcal{S}_{alg}^{(i)}) = \emptyset \mid \mathcal{B}_{alg}^{(i)}\right) \\ &= 1 - \left(1 - \frac{|\mathcal{S}^* \setminus \mathcal{S}_{alg}^{(i)}|}{|\mathcal{X}|}\right)^r \\ &= 1 - \left(1 - \frac{k-i}{m}\right)^r. \end{aligned} \quad (17)$$

Note that since $p_{alg}^{(i)} \leq 1$, it follows that

$$\prod_{i=0}^{k-1} p_{alg}^{(i)} \leq p_{alg}^{(k-1)} = 1 - \left(1 - \frac{1}{m}\right)^r. \quad (18)$$

Therefore, since $\max_i r_i < m - k$ and $\limsup_{m, k \rightarrow \infty} r_i/m = 0$, for all i , we can establish

$$\begin{aligned} \limsup_{m, k \rightarrow \infty} \Pr\left(\mathcal{S}_{alg}^{(k)} = \mathcal{S}^*\right) &\leq \limsup_{m, k \rightarrow \infty} \prod_{i=0}^{k-1} p_{alg}^{(i)} \leq \limsup_{m, k \rightarrow \infty} p_{alg}^{(k-1)} \\ &= \limsup_{m, k \rightarrow \infty} 1 - \left(1 - \frac{1}{m}\right)^r = 0. \end{aligned} \quad (19)$$

We next consider the case where $\mathcal{R}_{alg}^{(i)}$ is constructed by sampling the elements in $\mathcal{X} \setminus \mathcal{S}_{alg}^{(i)}$ without replacement. The probability of success in this case is higher than when sampling with replacement. Nevertheless, we can derive $p_{alg}^{(i)}$ according to

$$p_{alg}^{(i)} = 1 - \prod_{l=0}^{r_i-1} \left(1 - \frac{k-i}{m-l}\right). \quad (20)$$

To establish the asymptotic probability, we upperbound the success probability with $p_{alg}^{(k-1)}$. Once again, since $\limsup_{m, k \rightarrow \infty} r_i/m = 0$, for all i , one can observe that $\limsup_{m, k \rightarrow \infty} p_{alg}^{(k-1)}$, thereby proving the stated result. ■

Theorem 2 establishes an upper bound on the probability that a variant of GREEDY with a restricted search space constructed uniformly at random identifies \mathcal{S}^* exactly. The theorem states that as long as $\limsup_{m, k \rightarrow \infty} r_i/m = 0$, the asymptotic success probability is zero. To illustrate the implications of this theorem, we consider the scenario where $r_i = r$ is kept fixed, i.e., the size of the restricted search space is constant. This choice is

³Note that $\mathcal{S}_{alg}^{(k)}$ and \mathcal{S}^* are quantities that depend on m .

explored in [3]–[6] where bounds on the expected approximation factor of Algorithm 2 are derived with $r = \frac{m}{k} \log \frac{1}{\epsilon}$. Corollary 2.1 below derives bound on the success probability under this consideration.

Corollary 2.1: Under the assumptions of Theorem 2 and $r = \frac{m}{k} \log \frac{1}{\epsilon}$, consider a sequence of optimization problems $\mathcal{P}(m, k)$ as (2) where $m, k \rightarrow \infty$. The following claims hold:

- 1) If there exists $\alpha \in (0, 1)$ such that $r \leq k^{\alpha-1}m$ (equivalently, $\epsilon \geq \exp(-k^\alpha)$), then the probability that ALG succeeds on $\mathcal{P}(m, k)$ goes to zero, i.e.,

$$\limsup_{m, k \rightarrow \infty} \Pr \left(\mathcal{S}_{alg}^{(k)} = \mathcal{S}^* \right) = 0. \quad (21)$$

- 2) If there exists $\alpha_1 \in (0, 1)$ such that $r \leq \alpha_1 m$ (equivalently, $\epsilon \geq \exp(-\alpha_1 k)$), then the probability that ALG succeeds on $\mathcal{P}(m, k)$ satisfies

$$\limsup_{m, k \rightarrow \infty} \Pr \left(\mathcal{S}_{alg}^{(k)} = \mathcal{S}^* \right) \leq 1 - \exp(-\alpha_1) < 0.64. \quad (22)$$

Proof: For simplicity, we only consider the case of constructing \mathcal{R} via sampling with replacement as the case of sampling without replacement can be treated analogously.

First, consider the setting where $\epsilon \geq \exp(-k^\alpha)$, $0 < \alpha < 1$. In light of Theorem 2, to establish the first result it suffices to show that

$$\limsup_{m, k \rightarrow \infty} 1 - \left(1 - \frac{1}{m} \right)^r = 0. \quad (23)$$

Using Lemma 2 yields

$$\begin{aligned} 1 - \left(1 - \frac{1}{m} \right)^r &\leq 1 - \exp\left(-\frac{r}{m}\right) \left(1 - \frac{r}{m^2} \right) \\ &= 1 - \exp\left(\frac{\log \epsilon}{k}\right) \left(1 + \frac{\log \epsilon}{mk} \right) \\ &\leq 1 - \exp(-k^{\alpha-1}) \left(1 - \frac{k^{\alpha-1}}{m} \right), \end{aligned} \quad (24)$$

where for the last inequality we recall the assumption $\epsilon \geq \exp(-k^\alpha)$. The result is then established by noting $\limsup_{m, k \rightarrow \infty} \frac{k^{\alpha-1}}{m} = 0$, $\limsup_{k \rightarrow \infty} \exp(-k^{\alpha-1}) = 1$, and using the squeeze theorem.

Next, consider the second setting, i.e., $\epsilon \geq \exp(-\alpha_1 k)$, $0 < \alpha_1 < 1$. Following a similar approach, one obtains

$$\Pr \left(\mathcal{S}_{alg}^{(k)} = \mathcal{S}^* \right) \leq 1 - \exp(-\alpha_1) \left(1 - \frac{\alpha_1}{m} \right) := \delta_2. \quad (25)$$

Since the bound in (25) goes to $1 - \exp(-\alpha_1)$ as $m, k \rightarrow \infty$, it holds that $\delta_2 \leq 1 - \exp(-\alpha_1) < 1 - 1/e < 0.64$. ■

Corollary 2.1 establishes upper bounds on the probability that a variant of GREEDY with a restricted search space constructed uniformly at random identifies \mathcal{S}^* exactly in two scenarios: (i) If the size of the search space remains fixed in each iteration of ALG and the algorithm makes $\mathcal{O}(mk^\alpha)$ oracle calls for some $\alpha \in (0, 1)$, then the probability of the exact identification approaches zero as the problem dimension grows. (ii) If the size of the search space remains fixed in each iteration of ALG and strictly less than $\lfloor m \rfloor$, and the algorithm makes $\mathcal{O}(mk)$ oracle

calls, then although the probability of the exact identification may be nonzero, it is not asymptotically 1. Note that in many applications, including sparse reconstruction and sparse learning [23], [26], an arbitrarily high success probability is a condition required to establish any nontrivial sample complexity results, i.e., the minimum number of data points for successful recovery and prediction.

Remark 1: At the first glance one may suggest that if a fixed sample size is used throughout the GREEDY with random sampling, i.e., Algorithm 2, the probability that the algorithm recovers the optimal solution goes to zero when the size of the problem grows. However, this statement is not accurate. In fact, one can show

$$\prod_{i=0}^{k-1} p^{(i)} \geq \delta, \quad (26)$$

for a strictly positive δ . Further details regarding this remark are postponed to Appendix B.

A. Progressively-Increasing Random Sampling

The proofs of Theorem 2 and Corollary 2.1 reveal the underlying cause for the failure of Algorithm 2 to find the optimal solution: since the size of the search domain is fixed throughout the iterations, if Algorithm 2 successfully identifies elements from \mathcal{S}^* in earlier iterations, the chance of sampling new elements from \mathcal{S}^* significantly decreases in the subsequent ones. Therefore, success in earlier iterations increases the chance of failure to select new elements from \mathcal{S}^* in the subsequent ones. This phenomenon is not encountered in the GREEDY algorithm since in each iteration GREEDY considers all the elements in the ground set, including those in \mathcal{S}^* . Therefore, although in initial iterations Algorithm 2 may search smaller domains, if the goal is to identify exactly all the elements in \mathcal{S}^* , one should progressively increase the size of the search domain to improve the probability of success.

We thus propose a simple *increasing schedule* strategy of search spaces which grow to *ultimately reach size* m . In particular, we progressively increase the size of the search domain as the cardinality of the identified subset $\mathcal{S}^{(i)}$ grows. The proposed method, referred to as *Progressive Stochastic Greedy* (PSG), in the i th iteration samples $r_i = \frac{m}{k-i} \log \frac{1}{\epsilon}$ elements uniformly at random from \mathcal{X} to construct the search set $\mathcal{R}^{(i)}$. Following [3]–[6], we let ϵ , such that $e^{-k} \leq \epsilon \leq e^{-\frac{k}{m}}$, be a parameter which allows one to strike a desired balance between performance and complexity; the sampling may be with or without replacement. Note that since $r_i \leq m$ for all $i = 0, \dots, k-1$, for any iteration i such that $i \geq i_s := k - \log \frac{1}{\epsilon}$, we set r_i to its maximum value, m . Thus, this procedure can be interpreted as a hybrid scheme that provides a soft transition from a restricted search space with random sampling to GREEDY.

Remark 2: If a prior distribution is known on which elements constitute the optimal subset, that distribution can be leveraged to improve the performance of PSG. Leveraging this information has been pursued in the special case of the Bayesian sparse reconstruction problem (see, e.g., [43]).

B. Complexity Analysis of PSG

Recall that GREEDY and Algorithm 2 with $r_i = r = \frac{m}{k} \log \frac{1}{\epsilon}$, i.e., STOCHASTIC-GREEDY [3]–[6], require $\mathcal{O}(mk)$ and $\mathcal{O}(m \log \frac{1}{\epsilon})$ function evaluations, respectively. In the following, we express the complexity of PSG in terms of the total number of function evaluations throughout the iterations of the algorithm. Specifically, given our choice of r_i , we have that

$$\mathcal{O} \left(\sum_{i=0}^{k-\log \frac{1}{\epsilon}} \frac{m}{k-i} \log \frac{1}{\epsilon} + m \log \frac{1}{\epsilon} \right) \quad (27)$$

As an example, for $\epsilon \geq \exp(-k^\alpha)$ and $\epsilon \geq \exp(-\alpha k)$, $0 < \alpha < 1$, the complexity expression in (27) reduces to $\tilde{\mathcal{O}}(mk^\alpha)$ and $\mathcal{O}(mk)$, respectively, while for a constant ϵ the complexity reduces to $\mathcal{O}(m \log(\frac{1}{\epsilon}) \log k)$. Thus, the proposed method incurs at most a factor $\mathcal{O}(\log k)$ higher complexity than STOCHASTIC-GREEDY. As we show in the remainder of the next subsection, this relatively small increase in complexity is sufficient to satisfy a necessary condition for identifying \mathcal{S}^* .

C. Lower Bound on the Probability of Success

We now aim to determine whether PSG, i.e., Algorithm 2 with the proposed progressively-increasing schedule, can identify the optimal solution to $\mathcal{P}(m, k)$ as $m, k \rightarrow \infty$.

To answer the above question, one needs to establish a sufficient condition for the exact identification of \mathcal{S}^* or, equivalently, a lower bound on the probability of success of Algorithm 2. Recall from Lemma 1 that

$$\Pr \left(\mathcal{S}^{(k)} = \mathcal{S}^* \right) = \prod_{i=0}^{k-1} p^{(i)} \prod_{i=0}^{k-1} q^{(i)}. \quad (28)$$

Therefore, it suffices to derive nontrivial lower bounds on $\prod_{i=0}^{k-1} p^{(i)}$ and $\prod_{i=0}^{k-1} q^{(i)}$. In Theorem 3 below, we provide a lower bound on $\prod_{i=0}^{k-1} p^{(i)}$.

Theorem 3: Suppose the optimal solution to (2) is unique. Let $r_i = \min(\frac{m}{k-i} \log \frac{1}{\epsilon}, m)$, for all $i = 0, \dots, k-1$. Then,

$$p = \prod_{i=0}^{k-1} p^{(i)} \geq (1 - \epsilon)^{k - \log \frac{1}{\epsilon}}. \quad (29)$$

Furthermore, if $\epsilon = \frac{1}{k^\alpha}$ for some $\alpha > 1$, then $\limsup_{m, k \rightarrow \infty} p = 1$.

Proof: Note that since $r_i = m$ for all $i \geq k - \log \frac{1}{\epsilon}$, it follows that $p^{(i)} = 1$. Let us first consider the setting of sampling with replacement. There, it holds that

$$\begin{aligned} \prod_{i=0}^{k-1} p^{(i)} &= \prod_{i=0}^{k-\log \frac{1}{\epsilon}-1} \left(1 - \left(1 - \frac{k-i}{m} \right)^{r_i} \right) \\ &\geq \prod_{i=0}^{k-\log \frac{1}{\epsilon}-1} \left(1 - \exp \left(-r_i \frac{k-i}{m} \right) \right) \\ &= (1 - \epsilon)^{k - \log \frac{1}{\epsilon}}. \end{aligned} \quad (30)$$

Next, we consider the setting of sampling without replacement. For every $i < k - \log \frac{1}{\epsilon}$,

$$\begin{aligned} p^{(i)} &= 1 - \prod_{l=0}^{r_i-1} \left(1 - \frac{k-i}{m-l} \right) \\ &\stackrel{(a)}{\geq} 1 - \left(1 - \frac{k-i}{r_i} \sum_{l=0}^{r_i-1} \frac{1}{m-l} \right)^{r_i} \geq 1 - \left(1 - \frac{k-i}{m} \right)^{r_i} \\ &\stackrel{(b)}{\geq} 1 - \exp \left(-r_i \frac{k-i}{m} \right) = 1 - \epsilon, \end{aligned} \quad (31)$$

where (a) is obtained by using the inequality relating arithmetic and geometric means, and (b) is due to the fact that $(1+x)^y \leq e^{xy}$ for any real number $y \geq 1$. Therefore, just as in the case of sampling with replacement, (29) holds.

Finally, to establish the asymptotic result, using $(1+x)^y \leq e^{xy}$ we can show that

$$(1 - \epsilon)^{k - \log \frac{1}{\epsilon}} \geq 1 - k\epsilon. \quad (32)$$

Hence, if $\epsilon = \frac{1}{k^\alpha}$ for some $\alpha > 1$, $\limsup_{m, k \rightarrow \infty} p = 1$ by the squeeze theorem. \blacksquare

Note that, as we argued in Corollary 2.1, for the specific value of $\epsilon = \frac{1}{k^\alpha}$ the success probability of Algorithm 2 with $r_i = r = \frac{m}{k} \log \frac{1}{\epsilon}$ goes to zero. Additionally, it turns out that in the regime where $\epsilon \leq \frac{1}{k^\alpha}$, the proposed procedure requires $\mathcal{O}(m \log^2 k) = \tilde{\mathcal{O}}(m)$ function evaluations as opposed to $\mathcal{O}(m \log k)$ function evaluations of STOCHASTIC-GREEDY.

Since we have established a lower bound on $\prod_{i=0}^{k-1} p^{(i)}$ in (29), we need only derive a nontrivial lower bound on $q^{(i)}$ in order to show existence of a sufficient condition for the exact identification of \mathcal{S}^* , and establish a lower bound on the probability of success of Algorithm 2 with the proposed sampling strategy.

A lower bound on $q^{(i)}$ can be obtained by considering the conditions under which the largest marginal gain of elements in $\mathcal{R}^{(i)} \cap \mathcal{S}^*$ exceeds that in $\mathcal{R}^{(i)} \setminus \mathcal{S}^*$ for all $i = 0, \dots, k-1$, i.e.,

$$\max_{j \in \mathcal{R}^{(i)} \setminus \mathcal{S}^*} f_j(\mathcal{S}^{(i)}) < \max_{j \in \mathcal{R}^{(i)} \cap \mathcal{S}^*} f_j(\mathcal{S}^{(i)}), \quad (33)$$

with high probability. Doing so requires problem-dependent information and a general result cannot be derived. However, below we specify such a condition for the problem of sparse support selection.

In the task of sparse support selection [23], [26], we are given a linear measurement vector $\mathbf{y} = \mathbf{A}\mathbf{x} + \boldsymbol{\nu}$ where $\mathbf{x} \in \mathbb{R}^m$ is a k -sparse unknown vector, i.e., a vector with at most k non-zero components, $\mathbf{y} \in \mathbb{R}^n$ denotes the vector of measurements, $\mathbf{A} \in \mathbb{R}^{n \times m}$ is the coefficient matrix assumed to be full rank, and $\boldsymbol{\nu} \in \mathbb{R}^n$ denotes the additive measurement noise vector. The search for a sparse approximation of \mathbf{x} leads to the NP-hard cardinality-constrained least-squares problem

$$\underset{\mathbf{x}}{\text{minimize}} \quad \|\mathbf{y} - \mathbf{A}\mathbf{x}\|_2^2 \quad \text{subject to} \quad \|\mathbf{x}\|_0 \leq k, \quad (34)$$

which can be interpreted as an instance of (2) [13], [41].

Theorem 4 below states that PSG successfully recovers k -sparse \mathbf{x} with high probability as long as the number of measurements is linear in k (sparsity) and logarithmic in $\frac{m}{k}$, achieving

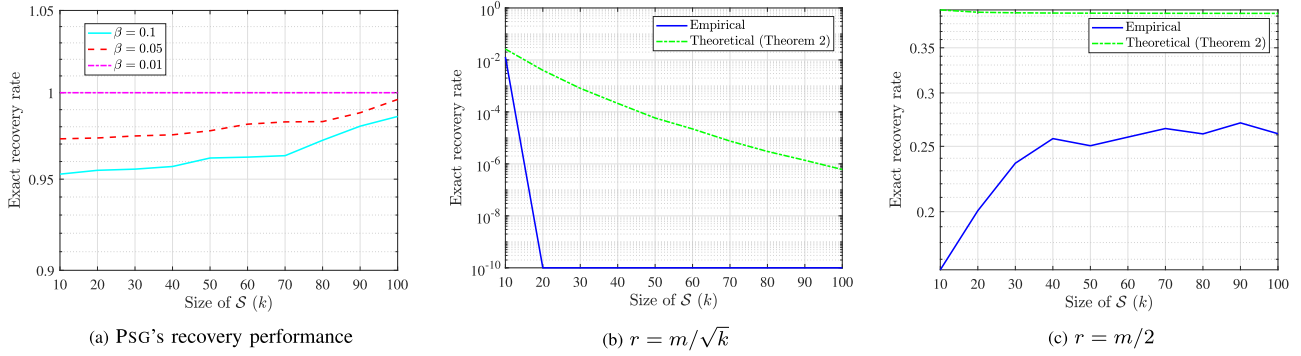


Fig. 1. Empirical evaluation of the theoretical bounds established by Theorem 2 and Corollary 2.1.

the optimal sample complexity established by Candes and Tao [26]. The proof is provided in the supplementary material.

Theorem 4: Let $\mathbf{x} \in \mathbb{R}^m$ be an arbitrary sparse vector with k non-zero entries and let $\mathbf{A} \in \mathbb{R}^{n \times m}$ denote a random matrix with entries drawn independently from $\mathcal{N}(0, 1/n)$. Given noiseless measurements $\mathbf{y} = \mathbf{A}\mathbf{x}$, PSG with parameter $e^{-k} \leq \epsilon \leq e^{-\frac{k}{m}}$ finds a solution that satisfies

$$\Pr(\mathcal{S}_{psg} = \mathcal{S}^*) \geq (1 - \epsilon)^{k - \log \frac{1}{\epsilon}} \left(1 - c_1 \left(\frac{m}{k} \right)^{c_2} \exp(-c_3 \frac{n}{k}) \right),$$

for some positive universal constants c_1, c_2 , and c_3 . Furthermore, assume that $m > k\sqrt{k}$ and

$$n \geq \max \left(\frac{6}{C_1} k \log \frac{m}{k \sqrt[6]{4\beta}}, C_2 k \right), \quad (35)$$

where $0 < \beta < 1$, and C_1 and C_2 are positive constants independent of β, n, m , and k . Then, PSG with parameter $\epsilon < \frac{\beta}{k}$ can exactly identify the optimal support subset \mathcal{S}^* with a probability of success exceeding $1 - 2\beta$.

D. Verifying the Theory

In this section, we verify our theoretical results by comparing them to the empirical ones obtained via Monte Carlo (MC) simulations. Specifically, we consider the task of sparse support selection.

We consider a setting with increasing support size k (varied from 10 to 100) and set the dimensions of the signal and the number of measurements to $m = 2k^{1.5}$ and $n = 6k \log(m/k \sqrt[6]{4\beta})$, respectively, for three different values of $\beta = 0.1, 0.05, 0.01$. In each trial, we select locations of the nonzero elements of \mathbf{x} uniformly at random and draw those elements from a normal distribution. Entries of the coefficient matrix \mathbf{A} are also generated randomly from $\mathcal{N}(0, \frac{1}{n})$. The results are averaged over 1000 Monte Carlo trials. Note that, as we show in the supplementary, in the above settings PSG is able to recover \mathbf{x} exactly.

First, we investigate the exact performance of PSG with the schedule

$$r_i = \frac{m}{k-i} \log \frac{1}{\epsilon}, \quad \epsilon = \frac{\beta}{k} \quad (36)$$

for $\beta = 0.1, 0.05, 0.01$, and show the results in Fig. 1(a). As can be seen there, the empirical exact recovery rate of PSG is very close to 1; this coincides with the theoretical lower bound of $1 - 2\beta$ established in the supplementary that builds upon the

insights of Theorem 2 (i.e, the achieved rate is 0.8, 0.9, 0.98 for $\beta = 0.1, 0.05, 0.01$, respectively).

Next, we empirically verify the results of Theorem 2 and Corollary 2.1 wherein we established an upper bound on the success probability of a variant of Algorithm 2, named ALG, with a restricted uniform search space. Fig. 1(b), (c) compares this theoretical result with the empirical success rate for $r = m/\sqrt{k}$ and $r = m/2$, which correspond to instances of the two settings considered in Corollary 2.1. Fig. 1(b) shows that for $r = m/\sqrt{k}$ the success rate goes to zero as k increases, as predicted by the first part of Corollary 2.1.⁴ In Fig. 1(c) we see that the success rate does not go to zero for $r = m/2$; moreover, it is always bounded by $1 - e^{-0.5} \approx 0.39$, as claimed by the second part of Corollary 2.1.

IV. ANALYSIS OF APPROXIMATION FACTOR

As argued in Section III, finding a lower bound on the success probability requires problem-dependent information and thus a general result cannot be derived. Therefore, in this section we focus on analyzing the performance of PSG from the perspective of establishing nontrivial worst-case approximation factors.

A. Expected Approximation Factor

We first establish bounds on the expected approximation factors of PSG. Then, assuming martingale structure of the marginal gains encountered in Algorithm 2, we establish a high probability result for the approximation factor and present it in the next subsection.

Proposition 2: Let \mathcal{S}_{psg} denote the random subset selected by PSG, and let c_f and e_f be the multiplicative and additive weak submodularity constants of the submodularity ratio of the set function objective in (2), respectively. Then

$$\mathbb{E}[f(\mathcal{S}_{psg})] \geq A_c f(\mathcal{S}^*), \quad \mathbb{E}[f(\mathcal{S}_{psg})] \geq A_e (f(\mathcal{S}^*) - (k-1)e_f), \quad (37)$$

where

$$A_c := 1 - \prod_{i=1}^{k - \log \frac{1}{\epsilon}} \left(1 - \frac{1 - \epsilon \frac{k}{k-i-1}}{kc} \right) \left(1 - \frac{1}{kc} \right)^{\log \frac{1}{\epsilon}}, \quad (38)$$

⁴Note that, in this setting, for $k \geq 20$ ALG failed in all of the trials; however, for illustration purposes (i.e., to be able to show the plot in the logarithmic scale) we set the success rate of ALG for $k \geq 20$ to 10^{-10} .

and

$$A_e := 1 - \prod_{i=1}^{k-\log \frac{1}{\epsilon}} \left(1 - \frac{1 - \epsilon^{\frac{k}{k-i-1}}}{k} \right) \left(1 - \frac{1}{k} \right)^{\log \frac{1}{\epsilon}}. \quad (39)$$

Proof: We first prove the result under a bounded c_f assumption. To simplify the notation, let $f_i := f(\mathcal{S}_{psg}^{(i)})$ and $f^* := f(\mathcal{S}^*)$.

Since $c_f \geq 1$, it holds that $\frac{1}{r}(1 + (r-1)c_f) \leq c_f$, for all r . Employing the first part of Lemma 1 with $\mathcal{S} = \mathcal{S}_{psg}^{(i)}$ and $\mathcal{T} = \mathcal{S}^* \cup \mathcal{S}_{psg}^{(i)}$, and recalling monotonicity of f , we obtain

$$\begin{aligned} f(\mathcal{S}^*) - f(\mathcal{S}_{psg}^{(i)}) &\leq f(\mathcal{S}^* \cup \mathcal{S}_{psg}^{(i)}) - f(\mathcal{S}_{psg}^{(i)}) \\ &\leq c_f \sum_{j \in \mathcal{S}^* \setminus \mathcal{S}_{psg}^{(i)}} f_j(\mathcal{S}_{psg}^{(i)}). \end{aligned} \quad (40)$$

We now proceed with the analysis by considering two distinct scenarios: the case where $r_i < m$ and hence the search space is strictly less than that of GREEDY, and the case where $r_i = m$. Note that the first iteration in which the latter condition holds is $i_s := k - \log \frac{1}{\epsilon}$.

For any $i < i_s$, we need to find the probability that the current search space $\mathcal{R}^{(i)}$ samples an element from the optimal subset. In light of the analysis in Section III, this probability for both sampling with and without replacement may be bounded below by

$$\begin{aligned} \Pr \left(\mathcal{R}^{(i)} \cap (\mathcal{S}^* \setminus \mathcal{S}^{(i)}) \neq \emptyset \right) &\geq 1 - e^{-\frac{r_i}{m} |\mathcal{S}^* \setminus \mathcal{S}^{(i)}|} \\ &\geq (1 - \epsilon^{\frac{k}{k-i}}) \left(\frac{|\mathcal{S}^* \setminus \mathcal{S}^{(i)}|}{k} \right). \end{aligned} \quad (41)$$

Given that for any $i < i_s$ PSG selects an element from $\mathcal{R}^{(i)}$ greedily, conditioned on $\mathcal{R}^{(i)} \cap (\mathcal{S}^* \setminus \mathcal{S}^{(i)}) \neq \emptyset$, the marginal gain of the selected value is larger than that of any element selected from $\mathcal{R}^{(i)} \cap (\mathcal{S}^* \setminus \mathcal{S}^{(i)})$ uniformly at random (an in turn from $\mathcal{S}^* \setminus \mathcal{S}^{(i)}$, given that the sampling is uniform). Therefore,

$$\mathbb{E} \left[f_{(i+1)}(\mathcal{S}_{psg}^{(i)}) | \mathcal{S}_{psg}^{(i)} \right] \geq \frac{1 - \epsilon^{\frac{k}{k-i}}}{|\mathcal{S}^* \setminus \mathcal{S}^{(i)}|} \sum_{j \in \mathcal{S}^* \setminus \mathcal{S}_{psg}^{(i)}} f_j(\mathcal{S}_{psg}^{(i)}). \quad (42)$$

Since $|\mathcal{S}^* \setminus \mathcal{S}^{(i)}| \leq k$, using (40) and taking the total expectation we obtain for all $i < i_s$

$$\mathbb{E} [(f_{i+1} - f_i)] \geq \frac{(1 - \epsilon^{\frac{k}{k-i}})}{kc_f} (f^* - \mathbb{E}[f_i]). \quad (43)$$

For $i \geq i_s$, the search space of PSG is equivalent to that of GREEDY. Therefore,

$$f_{i+1} - f_i \geq \frac{1}{kc_f} (f^* - \mathbb{E}[f_i]). \quad (44)$$

Let $\Delta_i = f^* - \mathbb{E}[f_i]$. Using induction,

$$\Delta_{k-1} \leq (1 - A_c) \Delta_0. \quad (45)$$

Rearranging and noting that $\Delta_0 = f^*$ since $f(\emptyset) = 0$ establishes the stated result.

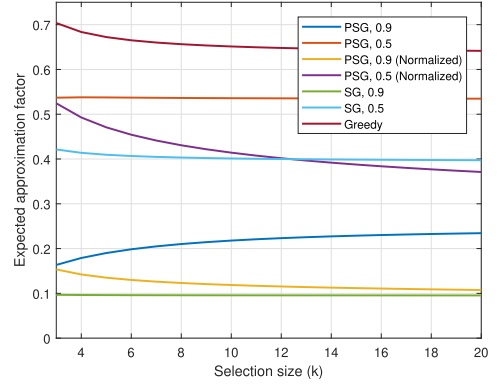


Fig. 2. Comparison of expected approximation factors of PSG, GREEDY, and STOCHASTIC-GREEDY for $\epsilon = 0.5, 0.9$ and different values of k . The normalized PSG variants have the same number of oracle calls as STOCHASTIC-GREEDY by employing $\epsilon' = \epsilon^{1/\log k}$.

The proof of the second part is analogous, except we leverage the second part of Lemma 1 and define $\Delta_i = f^* - \mathbb{E}[f_i] - (k-1)e_f$ to arrive at the recursion $\Delta_{k-1} \leq (1 - A_e)\Delta_0$, thereby completing the proof. ■

Note that, as expected, if the function f is submodular, i.e., $c_f = 1$ and $e_f = 0$, then $A_c = A_e$ and the expected approximation factors in (37) become identical.

It can be shown using elementary algebra that $A_c \geq 1 - e^{-\frac{1}{c_f}} - \frac{\epsilon}{c_f}$ and $A_e \geq 1 - \frac{1}{e} - \epsilon$. Such bounds are expected as the search space of PSG, i.e. r_i , can be lower bounded by that of the STOCHASTIC-GREEDY, $r = \frac{m}{k} \log(1/\epsilon)$. Therefore, the expected approximation factors of PSG are at least as large as those of STOCHASTIC-GREEDY. However, it is worth noting that these lower bounds are typically loose. To show this, we plot A_e and compare the results to the approximation factors of GREEDY and STOCHASTIC-GREEDY in Fig. 2 for two values of $\epsilon = 0.5, 0.9$. As the figure demonstrates, there is a large gap between the expected approximation factors of PSG and STOCHASTIC-GREEDY. Note that the improvement in the approximation factor only requires a marginal increase in the number of oracle calls [cf. (27)]. The figure also shows the approximation factor for a normalized variant of PSG. The normalized is performed such that, effectively, this variant of PSG employs $\epsilon' = \epsilon^{1/\log k}$. This choice ensures that this variant of PSG has the same number of oracle calls as STOCHASTIC-GREEDY with $\epsilon = 0.9, 0.5$. As the figure shows, as k grows the approximation factor of normalized PSG approaches that of STOCHASTIC-GREEDY. Overall, in scenarios where ϵ is large and k is small, which is typical in practice, PSG enjoys better guarantees.

B. High-Probability Bounds for the Approximation Factor

In Proposition 2, we established bounds on the expected approximation factor of the proposed scheme. Given that these results hold only on expectation, it is of interest to explore whether similar performance guarantees hold with high probability. The prior work [6] has shown a high probability bound for Algorithm 2 when $r_i = r = \frac{m}{k} \log \frac{1}{\epsilon}$. To derive this result, each step of Algorithm 2 is interpreted as an approximation of

the marginal gains of the selected elements using GREEDY,

$$f_{j_{(alg)}}(\mathcal{S}^{(i)}) = \eta^{(i)} f_{j_{(g)}}(\mathcal{S}^{(i)}), \quad (46)$$

where subscripts *alg* and *g* refer to the elements selected by Algorithm 2 and GREEDY, respectively, and $\{\eta^{(i)}\}_{i=1}^k$ is a collection of random variables in $(0,1]$.

The result of [6], however, relies on the assumption that $\{\eta^{(i)}\}_{i=1}^k$ are independent, which is not realistic since the element selected in each iteration depend on the selected subset so far. In Proposition 3, we relax that independence assumption and show that as long as $\{\eta^{(i)}\}_{i=1}^k$ form a martingale, one can derive a high-probability performance guarantee. Additionally, we derive high-probability bounds using both additive and multiplicative weak submodularity constants.

Proposition 3: Instate the notation and assumptions of Proposition 2. Let η be a martingale with $\mathbb{E}(\eta) \geq \mu$ satisfying the conditions of Theorem 5. Then, it holds with probability of at least $1 - \delta$ (for any $\delta > 0$) that

$$f(\mathcal{S}_{psg}) \geq A_c^{hp} f(\mathcal{S}^*), \quad f(\mathcal{S}_{psg}) \geq A_e^{hp} (f(\mathcal{S}^*) - (k-1)e_f), \quad (47)$$

where

$$A_c^{hp} := 1 - \exp\left(-\frac{1}{c_f} \left(\mu + \sqrt{\frac{\log \frac{1}{\delta}}{2k} + \frac{\log \frac{1}{\epsilon}}{k}}\right)\right), \quad (48)$$

and

$$A_e^{hp} := 1 - \exp\left(-\left(\mu + \sqrt{\frac{\log \frac{1}{\delta}}{2k} + \frac{\log \frac{1}{\epsilon}}{k}}\right)\right). \quad (49)$$

Proof: Let $f_i := f(\mathcal{S}_{psg}^{(i)})$ and $f^* := f(\mathcal{S}^*)$. Note that for any element j it holds that

$$f_j(\mathcal{S}_{psg}^{(i)}) \leq f_{(j)_{psg}}(\mathcal{S}_{psg}^{(i)}) \leq f_{(j)_g}(\mathcal{S}_{psg}^{(i)}). \quad (50)$$

Combining the above result with (40) and (46) we obtain

$$f_{i+1} - f_i \geq \frac{\eta^{(i+1)} \mathbb{I}(i < i_s) + \mathbb{I}(i \geq i_s)}{kc_f} (f^* - f_i). \quad (51)$$

Using an inductive argument similar to the one in the proof of Proposition 2, we obtain

$$f_{k-1} \geq f^* \left[1 - e^{-\frac{1}{kc_f} \sum_{i=1}^{i_s} \eta^{(i)}} \left(1 - \frac{1}{kc_f}\right)^{\log \frac{1}{\epsilon}} \right]. \quad (52)$$

Next, we use the result of Theorem 5 to bound the deviation of $\sum_{i=1}^{i_s} \eta^{(i)}$ from its mean with high probability. By Theorem 5 and the assumption that $\eta^{(i)} \in (0,1]$,

$$\Pr\left(\sum_{i=1}^{i_s} \eta^{(i)} - i_s \mu > \lambda\right) \leq \exp\left(-\frac{\lambda^2}{2i_s}\right). \quad (53)$$

Setting the right-hand side of the above inequality to δ yields that with probability exceeding $1 - \delta$

$$\sum_{i=1}^{i_s} \eta^{(i)} \leq i_s \mu + \sqrt{\frac{i_s}{2} \log \frac{1}{\delta}} \leq k\mu + \sqrt{\frac{k}{2} \log \frac{1}{\delta}}, \quad (54)$$

where we relied on the fact that $i_s \leq k$. Using this bound in (52) along with $(1+x)^y \leq e^{xy}$ for any real number $y \geq 1$ proves the first part of the proposition.

Leveraging the second part of Lemma 1 and repeating similar arguments furnishes the second result. ■

Proposition 3 establishes high probability bounds on the worst-case performance of the proposed scheme. For large k , the approximation factors may be approximated by $1 - e^{-\mu/c_f}$ and $1 - e^{-\mu}$ for multiplicative and additive weak submodularity constants, respectively. Therefore, if μ is near 1, as expected, the marginal gains of elements selected by PSG are relatively close to those that would have been selected by GREEDY, which in turn means we recover the guarantees of GREEDY in that regime.

V. APPLICATIONS TO SUBSET SELECTION

We demonstrate efficacy of the proposed algorithm in two applications, namely column subset selection for sparse subspace clustering and observation selection for target tracking.

A. Column Subset Selection for Subspace Clustering

Here we present results of an empirical evaluation of the proposed PSG scheme; specifically, the performance of PSG is compared to several baselines on the task of dimensionality reduction via column subset selection (CSS) [27] for sparse subspace clustering (SSC) [44]–[47].

The goal of CSS is to identify a subset \mathcal{S} , $|\mathcal{S}| = k$, of the set of m columns of a data matrix $\mathbf{A} \in \mathbb{R}^{n \times m}$ that best approximate the entire data matrix. Formally, the task of identifying \mathcal{S} can be cast as the optimization problem

$$\underset{\mathcal{S}}{\text{minimize}} \quad \|\mathbf{A} - \mathbf{P}_{\mathcal{S}} \mathbf{A}\|_F^2 \quad \text{s.t.} \quad |\mathcal{S}| = k, \quad (55)$$

where $\mathbf{P}_{\mathcal{S}} = \mathbf{A}_{\mathcal{S}} \mathbf{A}_{\mathcal{S}}^{\dagger}$ is the projection operator (specifically, projection onto the span of columns of $\mathbf{A}_{\mathcal{S}}$) and $\mathbf{A}_{\mathcal{S}}^{\dagger} = (\mathbf{A}_{\mathcal{S}}^{\top} \mathbf{A}_{\mathcal{S}})^{-1} \mathbf{A}_{\mathcal{S}}^{\top}$ denotes the Moore-Penrose pseudo-inverse of $\mathbf{A}_{\mathcal{S}}$. Since $\mathbf{A} = \mathbf{P}_{\mathcal{S}} \mathbf{A} + (\mathbf{I} - \mathbf{P}_{\mathcal{S}}) \mathbf{A}$ and $\|\mathbf{A}\|_F^2 = \|\mathbf{P}_{\mathcal{S}} \mathbf{A}\|_F^2 + \|(\mathbf{I} - \mathbf{P}_{\mathcal{S}}) \mathbf{A}\|_F^2$ by properties of projection matrices, (55) can equivalently be written as an instance of the weak submodular maximization task (2) [31], [48].

We aim to use Algorithm 2 as a CSS-based dimensionality reduction technique to reduce the cost of performing clustering via SSC [44], [45]. That is, using a lower dimensional data matrix \mathbf{A}_{S_g} obtained via CSS, we learn the representation matrix \mathbf{C} by solving

$$\underset{\mathbf{C}}{\text{minimize}} \quad \|\mathbf{A}_{S_g} - \mathbf{A}_{S_g} \mathbf{C}\|_F^2 + \lambda \|\mathbf{C}\|_1, \quad (56)$$

and then employ spectral clustering [49] on $\mathbf{W} = |\mathbf{C}| + |\mathbf{C}|^{\top}$ to segment the data points.

We consider the proposed PSG scheme with $\epsilon = 0.1$. We consider GREEDY, STOCHASTIC-GREEDY, and random column subset selection as the benchmarking schemes. Additionally, we use the best rank- k approximation of a matrix (i.e., top- k SVD) to serve as an upper bound on the achievable performance; note that this scheme explicitly minimizes the Frobenius reconstruction criterion. We compare performance of the above algorithms

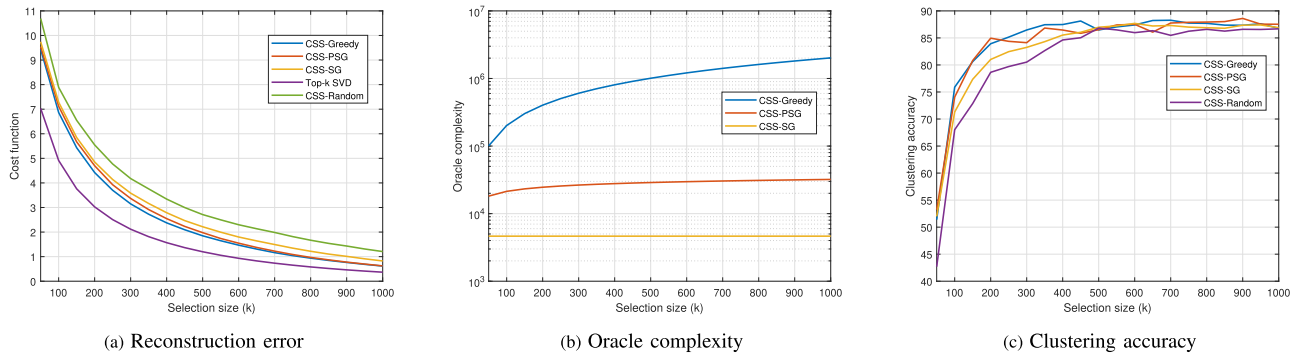


Fig. 3. Performance comparison of SSC with various CSS-based dimensionality reduction schemes on EYaleB dataset consisting of face images under 64 different illumination conditions.



Fig. 4. Face clustering: given images of multiple subjects, find images that belong to the same subject (examples from the EYaleB dataset [50]).

using the real EYaleB dataset [50] which contains frontal face images of 38 individuals under 64 different illumination conditions (see Fig. 4). There are $m = 2414$ columns (i.e., features) in this dataset; we select k out of $m = 2414$ columns, where k varies from 100 to 1000, and apply the SSC method of [45] to cluster the data points based on the selected features.

Fig. 3 shows the performance of various column subset selection schemes as well as the top- k SVD approach. In Fig. 3(a) we observe that the reconstruction errors of GREEDY and the proposed scheme are nearly identical, and that as we increase the number of selected columns the reconstruction error decreases; this is consistent with the fact that $f(\mathcal{S})$ is a monotone function. Fig. 3(b) shows a significant computational complexity improvement that the proposed scheme provides over the greedy CSS method. Since the complexity of Algorithm 2 increases logarithmically in k , the cost of selecting more columns is relatively small compared to the greedy approach. While we observe $\epsilon = 0.1$ achieves a good tradeoff between computational costs and performance, depending on the amount of data available, the value of ϵ can be tuned using cross-validation. Finally, in Fig. 3(c) we compare the clustering accuracy of SSC applied to a subset of features selected by different schemes. As the figure shows, clustering performance of SSC combined with the proposed CSS method is nearly identical to that of the conventional greedy approach; moreover, both achieve superior accuracy compared to the scheme that randomly select subsets of columns.

B. Observation Selection in Target Tracking

We consider a realistic multi-object tracking application by an autonomous car, referred to as the ego vehicle, that moves on a highway. Perception plays a major role in the planning subsystem of an autonomous vehicle; accurately identifying and

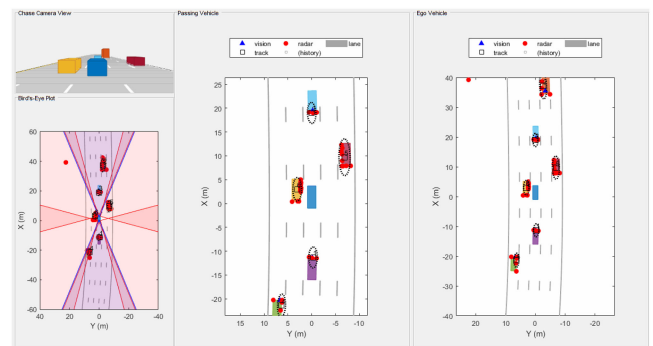




Fig. 5. The object tracking scenario: an ego vehicle (dark blue) moves on a highway with five lanes and aims to identify and track the vehicles in its surroundings using six radar and two camera sensors.

tracking the critical objects in its surrounding, e.g., other vehicles on the highway, are critical in the design of the perception subsystem [51]. Given the typically large amount of data gathered by the vehicle's sensors, observation selection techniques may be employed to reduce the computational burden of tracking algorithms and improve their runtime [6], [52]. Motivated by this argument, we resort to performing observation selection via the proposed scheme.

We consider a scenario with an ego vehicle that moves on a stretch of 500 meters of a typical highway road with five lanes, along with six other vehicles (see Fig. 5 and the supplementary gif file). All the vehicles move with the speed of 25 km/s, except for the passing vehicle which has a speed of 35 km/s. The ego vehicle has six radar sensors and two vision sensors covering the 360 degrees field of view. The sensors have some overlap and some coverage gap. The ego vehicle is equipped with a long-range radar sensor and a vision sensor on both



(a) Tracking error



(b) Runtime

Fig. 6. Performance comparison of the proposed observation selection algorithm with the greedy approach, and an approach that utilizes all of the gathered observations. The proposed scheme results in nearly the same OSPA while significantly reducing the running time of tracking.

the front and the back of the vehicle. Each side of the vehicle has two short-range radar sensors, each covering 90 degrees. One sensor on each side covers the area from the middle of the vehicle to the back. The other sensor on each side covers the area from the middle of the vehicle forward. We set the properties of the sensors such that they are in agreement with the realistic characteristics of radar and camera sensors [53]. The sensors gather noisy observations about the location and velocity of the surrounding vehicles. After selecting the observations via a subset selection algorithm, the selected observations are used in a GGIW-PHD extended object tracker that relies on a rectangular target model for the surrounding vehicles [54]–[56]. We assess the performance of the observation selection and tracking algorithms based on the optimal sub-pattern assignment (OSPA) metric [57], [58] as well as the time it takes to produce the estimated locations.

Fig. 6 shows the performance comparison of the proposed scheme with greedy observation selection that selects 50% of the observations. We also provide a comparison with a method that uses all of the gathered observations in the tracking algorithm. As Fig. 6(a) shows, all schemes achieve similar tracking accuracy. Note that OSPA increases from time $t = 50$ to $t = 75$ as in this interval the passing vehicle gets farther from the ego vehicle,

but it is still considered as a tracked object in the OSPA metric. However, after $t = 75$, the passing vehicle is not in the line of sight of the ego car anymore, resulting in the OSPA metric to improve.

Fig. 6(b) depicts the runtime comparison and demonstrates that using the proposed scheme to perform observation selection significantly reduces the cost of object tracking. Note that in the interval (35,50), the passing vehicle is in the proximity of the ego vehicle. Therefore, the mounted sensors produce significantly more observations, thereby increasing the runtime of the benchmarking schemes. However, thanks to the progressive random sampling strategy, the runtime of the proposed scheme does not increase significantly in this interval.

VI. CONCLUSION

In this paper, we studied the problem of large-scale monotone weak submodular maximization that comes up in many modern signal processing and machine learning applications including sparse reconstruction, dimensionality reduction, observation gathering, and sensor selection. Motivated by the desire to reduce complexity of the celebrated greedy scheme, we theoretically studied fundamental performance limits of restricting the size of the greedy search space by means of uniform sampling strategies. We first studied the asymptotic probability of successfully identifying the optimal subset. Our analysis revealed that many of the standard practices that rely on fixed sampling sizes lead to a success probability that is asymptotically zero. We showed that an increasing schedule of the search space size satisfies a necessary condition for the exact identification of the optimal subset in large-scale problems. Following this insight, we proposed a progressive stochastic greedy algorithm and demonstrated its efficacy in the applications to sparse subset selection, dimensionality reduction, and extended object tracking. We further established strong guarantees, both on expectations and with high probability, on the approximation factors of the proposed algorithm.

Our established framework gives rise to interesting open problems. First, the proposed analysis can be employed to study exact identification conditions of PSG in other classes of weak submodular maximization problems, i.e., settings beyond the sparse support selection task considered in this paper. For instance, exact identification conditions of GREEDY for the task of observation selection were recently considered by [59]. Utilizing our framework, similar results may be established for PSG. Furthermore, as we argued, the intersection of search space of PSG and new elements from the optimal subset is nonempty with high probability for all iterations. Since GREEDY considers all the elements of the ground set in each iteration, the intersection of the search space of GREEDY and new elements from the optimal subset is always nonempty. Therefore, it is reasonable to ask whether the set of conditions under which GREEDY and PSG exactly identify the optimal subset of a given weak submodular optimization problem \mathcal{P} is the same. Indeed, this is the case for the problem of sparse support selection. Finally, studying general weak submodular optimization problems remains of interest.

APPENDIX A
AUXILIARY LEMMA

Lemma 2: For every $|a| \leq 1$ and $b \geq 1$ it holds that $(1 + a)^b \geq e^{ab}(1 - a^2b)$.

Proof: Let $x = ab$, $|x| \leq b$. Consider $g(x) = e^{-x}(1 + \frac{x}{b})^b - (1 - \frac{x^2}{b})$. At $x = 0$, both $g(x)$ and $f'(x)$ are zero. If $f'(x) = 0$ for any other x in the interval, for such x we have

$$e^{-x} \left(1 + \frac{x}{b}\right)^b = 2 + \frac{2x}{b}.$$

Therefore, for such x

$$g(x) = \frac{(x+1)^2}{b} + 1 - \frac{1}{b} > 0.$$

Furthermore, since $g(b) > 0$ for all b while $g(-b) > 0$ for $b > 1$ and $g(-b) = 0$ for $b = 1$, all other points we must have $g(x) > 0$. ■

APPENDIX B
ESTABLISHING (26)

Let us consider the second case of Corollary 2.1, i.e., there exists $\alpha_1 \in (0, 1)$ such that $r = \alpha_1 m$ (equivalently, $\epsilon = \exp(-\alpha_1 k)$). In Corollary 2.1 we established an upperbound for the success probability which shows as the size of the problem grows the upper bound is strictly less than 0.64. However, this does not mean the probability approaches zero. To see this, upon recalling (11) we may write

$$\begin{aligned} \prod_{i=0}^{k-1} p^{(i)} &\geq \prod_{i=0}^{k-1} \left(1 - \exp\left(-r \frac{k-i}{m}\right)\right) \\ &= \prod_{i=0}^{k-1} (1 - \exp(-\alpha_1(k-i))) \\ &= \prod_{n=1}^k (1 - \exp(-\alpha_1 n)). \end{aligned} \tag{57}$$

A crude lowerbound can be easily achieved by replacing the product terms with the smallest term, i.e., $1 - \exp(-\alpha_1)$. Such lowerbound – as intuition dictates – converges to zero as $k \rightarrow \infty$. However, this is not a tight bound. Note that for every $0 < \alpha_1 < 1$, there exists a positive integer $q(\alpha_1)$ such that $\exp(\alpha_1 n) \geq n^2$ for all $n \geq q(\alpha_1)$. For instance, if $\alpha_1 \geq \log(2) \approx 0.6931$, then $q(\alpha_1) \leq 4$. Hence,

$$\prod_{n=q(\alpha_1)}^{\infty} \left(1 - \frac{1}{n^2}\right) \leq \prod_{n=q(\alpha_1)}^{\infty} \left(1 - \frac{1}{\exp(\alpha_1 n)}\right). \tag{58}$$

It then follows

$$\begin{aligned} \lim_{k \rightarrow \infty} \prod_{i=0}^{k-1} p^{(i)} &\geq \prod_{n=1}^{\infty} \left(1 - \frac{1}{\exp(\alpha_1 n)}\right) \\ &\geq \left(1 - \frac{1}{\exp(\alpha_1)}\right) \prod_{n=2}^{q(\alpha_1)-1} \frac{\left(1 - \frac{1}{\exp(\alpha_1 n)}\right)}{\left(1 - \frac{1}{n^2}\right)} \prod_{n=2}^{\infty} \left(1 - \frac{1}{n^2}\right) \end{aligned}$$

$$= \frac{1}{2} \left(1 - \frac{1}{\exp(\alpha_1)}\right) \prod_{n=2}^{q(\alpha_1)-1} \frac{\left(1 - \frac{1}{\exp(\alpha_1 n)}\right)}{\left(1 - \frac{1}{n^2}\right)} := \delta > 0. \tag{59}$$

For instance, if $\alpha_1 \geq \log(2)$, then $\delta \geq 0.2461$. Therefore, despite the intuition, the probability approaches a strictly nonzero value as the size of the problem grows.

The above argument does not take $\lim_{k \rightarrow \infty} \prod_{i=0}^{k-1} q^{(i)}$ into account and this probability may approach zero; however, its analysis does not seem trivial.

APPENDIX C
AZUMA’S INEQUALITY [60]

Theorem 5: Let X be a martingale associated with a filtration \mathcal{F} and a sequence of random variables X_0, X_1, \dots, X_n satisfying $X_i = \mathbb{E}[X|\mathcal{F}_i]$ and, in particular, $X_0 = \mathbb{E}[X]$ and $X_n = X$. If $|X_i - X_{i-1}| \leq c_i$, then,

$$\Pr(X - \mathbb{E}[X] > \lambda) \leq \exp\left(-\frac{\lambda^2}{2 \sum_{i=1}^n c_i^2}\right). \tag{60}$$

REFERENCES

- [1] A. Hashemi, H. Vikalo, and G. de Veciana, “On the performance-complexity tradeoff in stochastic greedy weak submodular optimization,” in *Proc. IEEE Int. Conf. Acoust., Speech Signal Process.*, 2021, pp. 3540–3544.
- [2] G. L. Nemhauser, L. A. Wolsey, and M. L. Fisher, “An analysis of approximations for maximizing submodular set functions—I,” *Math. Program.*, vol. 14, no. 1, pp. 265–294, 1978.
- [3] B. Mirzasoleiman, A. Badanidiyuru, A. Karbasi, J. Vondrak, and A. Krause, “Lazier than lazy greedy,” in *Proc. 29th AAAI Conf. Artif. Intell.*, 2015, pp. 1812–1818.
- [4] E. R. Elenberg, R. Khanna, A. G. Dimakis, and S. Negahban, “Restricted strong convexity implies weak submodularity,” *Ann. Statist.*, vol. 46, no. 6B, pp. 3539–3568, 2018.
- [5] A. Hashemi, M. Ghasemi, H. Vikalo, and U. Topcu, “A randomized greedy algorithm for near-optimal sensor scheduling in large-scale sensor networks,” in *Proc. Amer. Control Conf.*, 2018, pp. 1027–1032.
- [6] A. Hashemi, M. Ghasemi, H. Vikalo, and U. Topcu, “Randomized greedy sensor selection: Leveraging weak submodularity,” *IEEE Trans. Autom. Control*, vol. 66, no. 1, pp. 199–212, Jan. 2021.
- [7] D. P. Williamson and D. B. Shmoys, *The Design of Approximation Algorithms*. Cambridge, U.K.: Cambridge Univ. Press, 2011.
- [8] D. Kempe, J. Kleinberg, and É. Tardos, “Maximizing the spread of influence through a social network,” in *Proc. 9th ACM SIGKDD Int. Conf. Knowl. Discov. Data Mining*, 2003, pp. 137–146.
- [9] A. Krause, A. Singh, and C. Guestrin, “Near-optimal sensor placements in Gaussian processes: Theory, efficient algorithms and empirical studies,” *J. Mach. Learn. Res.*, vol. 9, pp. 235–284, 2008.
- [10] A. Guillory and J. Bilmes, “Active semi-supervised learning using submodular functions,” in *Proc. Twenty-Seventh Conf. Uncertainty Artif. Intell.*, 2011.
- [11] J. A. Tropp and A. C. Gilbert, “Signal recovery from random measurements via orthogonal matching pursuit,” *IEEE Trans. Inf. Theory*, vol. 53, no. 12, pp. 4655–4666, Dec. 2007.
- [12] S. Joshi and S. Boyd, “Sensor selection via convex optimization,” *IEEE Trans. Signal Process.*, vol. 57, no. 2, pp. 451–462, Feb. 2009.
- [13] A. Das and D. Kempe, “Submodular meets spectral: Greedy algorithms for subset selection, sparse approximation and dictionary selection,” in *Proc. Int. Conf. Mach. Learn.*, 2011, pp. 1057–1064.
- [14] B. Mirzasoleiman, A. Karbasi, R. Sarkar, and A. Krause, “Distributed submodular maximization: Identifying representative elements in massive data,” in *Proc. Adv. Neural Inf. Process. Syst.*, 2013, pp. 2049–2057.
- [15] E. Balkanski and Y. Singer, “The adaptive complexity of maximizing a submodular function,” in *Proc. 50th Annu. ACM SIGACT Symp. Theory Comput.*, 2018, pp. 1138–1151.

- [16] A. Ene and H. L. Nguyen, "Submodular maximization with nearly-optimal approximation and adaptivity in nearly-linear time," in *Proc. 30th Annu. ACM-SIAM Symp. Discrete Algorithms*, 2019, pp. 274–282.
- [17] M. Fahrback, V. Mirrokni, and M. Zadimoghaddam, "Submodular maximization with optimal approximation, adaptivity and query complexity," in *Proc. 2019 Annu. ACM-SIAM Symp. Discrete Algorithms*, 2019, pp. 255–273.
- [18] M. Minoux, "Accelerated greedy algorithms for maximizing submodular set functions," in *Proc. Optim. Techn.*, 1978, pp. 234–243.
- [19] A. Badanidiyuru and J. Vondrák, "Fast algorithms for maximizing submodular functions," in *Proc. 25th Annu. ACM-SIAM Symp. Discrete Algorithms*, 2014, pp. 1497–1514.
- [20] T. Zhang, "Sparse recovery with orthogonal matching pursuit under RIP," *IEEE Trans. Inf. Theory*, vol. 57, no. 9, pp. 6215–6221, Sep. 2011.
- [21] M. A. Davenport and M. B. Wakin, "Analysis of orthogonal matching pursuit using the restricted isometry property," *IEEE Trans. Inf. Theory*, vol. 56, no. 9, pp. 4395–4401, Sep. 2010.
- [22] Q. Mo and Y. Shen, "A remark on the restricted isometry property in orthogonal matching pursuit," *IEEE Trans. Inf. Theory*, vol. 58, no. 6, pp. 3654–3656, Jun. 2012.
- [23] J. A. Tropp, "Greedy is good: Algorithmic results for sparse approximation," *IEEE Trans. Inf. Theory*, vol. 50, no. 10, pp. 2231–2242, Oct. 2004.
- [24] T. T. Cai and L. Wang, "Orthogonal matching pursuit for sparse signal recovery with noise," *IEEE Trans. Inf. Theory*, vol. 57, no. 7, pp. 4680–4688, Jul. 2011.
- [25] T. Zhang, "On the consistency of feature selection using greedy least squares regression," *J. Mach. Learn. Res.*, vol. 10, pp. 555–568, Mar. 2009.
- [26] E. J. Candes and T. Tao, "Decoding by linear programming," *IEEE Trans. Inf. Theory*, vol. 51, no. 12, pp. 4203–4215, Dec. 2005.
- [27] J. A. Tropp, "Column subset selection, matrix factorization, and eigenvalue optimization," in *Proc. 20th Annu. ACM-SIAM Symp. Discrete Algorithms*, 2009, pp. 978–986.
- [28] V. Guruswami and A. K. Sinop, "Optimal column-based low-rank matrix reconstruction," in *Proc. 23rd Annu. ACM-SIAM Symp. Discrete Algorithms*, 2012, pp. 1207–1214.
- [29] C. Boutsidis, P. Drineas, and M. Magdon-Ismael, "Near-optimal column-based matrix reconstruction," *SIAM J. Comput.*, vol. 43, no. 2, pp. 687–717, 2014.
- [30] A. K. Farahat, A. Elgohary, A. Ghodsi, and M. S. Kamel, "Distributed column subset selection on mapreduce," in *Proc. 13th IEEE Int. Conf. Data Mining*, 2013, pp. 171–180.
- [31] A. K. Farahat, A. Elgohary, A. Ghodsi, and M. S. Kamel, "Greedy column subset selection for large-scale data sets," *Knowl. Inf. Syst.*, vol. 45, no. 1, pp. 1–34, 2015.
- [32] A. Civril and M. Magdon-Ismael, "Column subset selection via sparse approximation of SVD," *Theor. Comput. Sci.*, vol. 421, pp. 1–14, 2012.
- [33] J. Altschuler, A. Bhaskara, G. Fu, V. Mirrokni, A. Rostamizadeh, and M. Zadimoghaddam, "Greedy column subset selection: New bounds and distributed algorithms," in *Proc. 33rd Int. Conf. Mach. Learn.*, 2016, pp. 2539–2548.
- [34] A. Krause and D. Golovin, "Submodular function maximization," in *Tractability: Practical Approaches to Hard Problems*, Cambridge, U.K.: Cambridge Univ. Press, 2014, pp. 71–104.
- [35] U. Feige, "A threshold of $\ln n$ for approximating set cover," *J. ACM*, vol. 45, no. 4, pp. 634–652, Jul. 1998.
- [36] H. Zhang and Y. Vorobeychik, "Submodular optimization with routing constraints," in *Proc. 30th AAAI Conf. Artif. Intell.*, 2016, pp. 819–825.
- [37] L. Chamon and A. Ribeiro, "Approximate supermodularity bounds for experimental design," in *Proc. Adv. Neural Inf. Process. Syst.*, 2017, pp. 5409–5418.
- [38] A. Hashemi, M. Ghasemi, and H. Vikalo, "Submodular observation selection and information gathering for quadratic models," in *Proc. 36th Int. Conf. Mach. Learn.*, 2019, pp. 2653–2662.
- [39] T. Horel and Y. Singer, "Maximization of approximately submodular functions," in *Proc. Adv. Neural Inf. Process. Syst.*, 2016, pp. 3045–3053.
- [40] A. A. Bian, J. M. Buhmann, A. Krause, and S. Tschitschek, "Guarantees for greedy maximization of non-submodular functions with applications," in *Proc. Int. Conf. Mach. Learn.*, 2017, pp. 498–507.
- [41] R. Khanna, E. Elenberg, A. Dimakis, S. Negahban, and J. Ghosh, "Scalable greedy feature selection via weak submodularity," in *Proc. Int. Conf. Artif. Intell. Statist.*, 2017, pp. 1560–1568.
- [42] M. Ghasemi, A. Hashemi, U. Topcu, and H. Vikalo, "On submodularity of quadratic observation selection in constrained networked sensing systems," in *Proc. Amer. Control Conf.*, 2019, pp. 4671–4676.
- [43] N. Lee, "Map support detection for greedy sparse signal recovery algorithms in compressive sensing," *IEEE Trans. Signal Process.*, vol. 64, no. 19, pp. 4987–4999, Oct. 2016.
- [44] E. Elhamifar and R. Vidal, "Sparse subspace clustering," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2009, pp. 2790–2797.
- [45] C. You, D. Robinson, and R. Vidal, "Scalable sparse subspace clustering by orthogonal matching pursuit," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2016, pp. 3918–3927.
- [46] A. Hashemi and H. Vikalo, "Accelerated orthogonal least-squares for large-scale sparse reconstruction," *Digit. Signal Process.*, vol. 82, pp. 91–105, 2018.
- [47] A. Hashemi and H. Vikalo, "Evolutionary self-expressive models for subspace clustering," *IEEE J. Sel. Topics Signal Process.*, vol. 12, no. 6, pp. 1534–1546, Dec. 2018.
- [48] A. Bhaskara, A. Rostamizadeh, J. Altschuler, M. Zadimoghaddam, T. Fu, and V. Mirrokni, "Greedy column subset selection: New bounds and distributed algorithms," in *Proc. 33rd Int. Conf. Mach. Learn.*, 2016, pp. 2539–2548.
- [49] A. Y. Ng, M. I. Jordan, and Y. Weiss, "On spectral clustering: Analysis and an algorithm," in *Proc. Adv. Neural Inf. Process. Syst.*, 2002, pp. 849–856.
- [50] A. S. Georghiades, P. N. Belhumeur, and D. J. Kriegman, "From few to many: Illumination cone models for face recognition under variable lighting and pose," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 23, no. 6, pp. 643–660, Jun. 2001.
- [51] M. Ghasemi, E. Bulgur, and U. Topcu, "Task-oriented active perception and planning in environments with partially known semantics," in *Proc. Int. Conf. Mach. Learn.*, 2020, pp. 3484–3493.
- [52] M. Ghasemi and U. Topcu, "Online active perception for partially observable Markov decision processes with limited budget," in *Proc. 58th IEEE Conf. Decis. Control*, 2019, pp. 6169–6174.
- [53] MATLAB, "Automated driving toolbox: Design, simulate, and test adas and autonomous driving systems," 2021. [Online]. Available: https://www.mathworks.com/help/driving/index.html?s_tid=CRUX_lftnav
- [54] K. Granström, C. Lundquist, and U. Orguner, "Tracking rectangular and elliptical extended targets using laser measurements," in *Proc. 14th Int. Conf. Inf. Fusion*, 2011, pp. 1–8.
- [55] K. Granström, C. Lundquist, F. Gustafsson, and U. Orguner, "On extended target tracking using PhD filters," *Unpublished technical report*, 2012. [Online]. Available: <http://users.isy.liu.se/rtr/karl/publications/GranstromLGO2012.pdf>
- [56] K. Granstrom, M. Baum, and S. Reuter, "Extended object tracking: Introduction, overview and applications," *J. Adv. Inf. Fusion*, vol. 12, no. 2, Dec. 2017.
- [57] D. Schuhmacher, B.-T. Vo, and B.-N. Vo, "A consistent metric for performance evaluation of multi-object filters," *IEEE Trans. Signal Process.*, vol. 56, no. 8, pp. 3447–3457, Aug. 2008.
- [58] B. Ristic, B.-N. Vo, D. Clark, and B.-T. Vo, "A metric for performance evaluation of multi-target tracking algorithms," *IEEE Trans. Signal Process.*, vol. 59, no. 7, pp. 3452–3457, Jul. 2011.
- [59] D. Sharma, A. Kapoor, and A. Deshpande, "On greedy maximization of entropy," in *Proc. Int. Conf. Mach. Learn.*, 2015, pp. 1330–1338.
- [60] F. Chung and L. Lu, "Concentration inequalities and martingale inequalities: A survey," *Internet Math.*, vol. 3, no. 1, pp. 79–127, 2006.



Abolfazl Hashemi (Member, IEEE) received the B.Sc. degree in electrical engineering from the Sharif University of Technology, Tehran, Iran, in July 2014, and the M.S.E. and Ph.D. degrees in electrical and computer engineering from the University of Texas at Austin, Austin, TX, USA, in May 2016 and August 2020, respectively. From August 2020 to August 2021, he was a Postdoctoral Scholar with the Oden Institute for Computational Engineering and Sciences, University of Texas at Austin. Since August 2021, he has been an Assistant Professor with the Elmore Family School of Electrical and Computer Engineering, Purdue University, West Lafayette, IN, USA. His research interests include distributed machine learning, optimization, signal processing, and control. Abolfazl was the 2019 Schmidt Science Fellows Award nominee from UT Austin, and was the recipient of the Iranian National Elite Foundation Fellowship and a Best Student Paper Award finalist at the 2018 American Control Conference.



Haris Vikalo (Senior Member, IEEE) received the B.S. degree in electrical engineering from the University of Zagreb, Zagreb, Croatia, in 1995, the M.S. degree in electrical engineering from Lehigh University, Bethlehem, PA, USA, in 1997, and the Ph.D. degree in electrical engineering from Stanford University, Stanford, CA, USA, in 2003. He held a short-term appointment with Bell Laboratories, Murray Hill, NJ, USA, in 1999. From January 2003 to July 2003, he was a Postdoctoral Researcher, and from July 2003 to August 2007, he was an Associate Scientist with

the California Institute of Technology, Pasadena, CA, USA. Since September 2007, he has been with the Department of Electrical and Computer Engineering, The University of Texas at Austin, Austin, TX, USA, where he is currently an Associate Professor. His research interests include signal processing, machine learning, communications, and bioinformatics. He was the recipient of the 2009 National Science Foundation Career Award.



Gustavo de Veciana (Fellow, IEEE) received the Ph.D. degree in electrical engineering from the University of California, Berkeley, Berkeley, CA, USA, in 1993. He is currently a Professor and the Associate Chair of the Department of Electrical and Computer Engineering and was the Cockrell Family Regents Chair in engineering with The University of Texas at Austin, Austin, TX, USA. He was the Director and Associate Director of the Wireless Networking and Communications Group (WNCG) from 2003–2007.

His research interests include design, analysis and control networks, information theory and applied probability. His current research interests include measurement, modeling and performance evaluation, wireless and sensor networks; architectures and algorithms to design reliable computing and networked systems. He is currently the Editor at Large for the IEEE/ACM TRANSACTIONS ON NETWORKING. He was the recipient of an NSF CAREER Award 1996, and a co-recipient of seven best paper awards including the 2021 IEEE Communication Society W. Bennett Prize. In 2009, he was designated IEEE Fellow for his contributions to the analysis and design of communication networks. He currently serves on the board of trustees of IMDEA Networks Madrid.