AFRL-AFOSR-VA-TR-2024-0217

Data Acquisition in Dynamic Environments: A Submodular Perspective

HASSANI, HAMED
TRUSTEES OF THE UNIVERSITY OF PENNSYLVANIA
3451 WALNUT ST
PHILADELPHIA, PA,
US

05/02/2024
Final Technical Report

Air Force Research Laboratory
Air Force Office of Scientific Research
Arlington, Virginia 22203
Air Force Materiel Command

REPORT DOCUMENTATION PAGE

PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ORGANIZATION.

1. REPORT DATE	2. REPORT TYPE	3. DATES COVERED	
20240502	Final	START DATE 20200601	END DATE 20230531

4. TITLE AND SUBTITLE
Data Acquisition in Dynamic Environments: A Submodular Perspective

5a. CONTRACT NUMBER	5b. GRANT NUMBER	5c. PROGRAM ELEMENT NUMBER
	FA9550-20-1-0111	61102F
5d. PROJECT NUMBER	**5e. TASK NUMBER**	**5f. WORK UNIT NUMBER**

6. AUTHOR(S)
HAMED HASSANI

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)
TRUSTEES OF THE UNIVERSITY OF PENNSYLVANIA
3451 WALNUT ST
PHILADELPHIA, PA
US

8. PERFORMING ORGANIZATION REPORT NUMBER

9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)
Air Force Office of Scientific Research
875 N. Randolph St. Room 3112
Arlington, VA 22203

10. SPONSOR/MONITOR'S ACRONYM(S)
AFRL/AFOSR RTA2

11. SPONSOR/MONITOR'S REPORT NUMBER(S)
AFRL-AFOSR-VA-TR-2024-0217

12. DISTRIBUTION/AVAILABILITY STATEMENT
A Distribution Unlimited: PB Public Release

13. SUPPLEMENTARY NOTES

14. ABSTRACT
While the recent advances in Artificial Intelligence have mainly relied on the availability of a wealth of centralized data, a fundamental challenge in many DoD-relevant applications is to acquire high-quality data at minimal cost. These applications range from robotic sensing and autonomous planning to experimental design and active learning; furthermore, such
applications often take place in unknown or even adversarial environments, in which data is highly limited and precious. More specifically, each observation may significantly impact our ability to learn and operate in unknown and dynamic environments. Moreover, dealing with complex real-world environments requires a paradigm shift from the existing static, modelaware
data acquisition approaches to methods that learn adaptively and are robust against imperfect, stochastic and evolving knowledge. Whether we select a bunch of sensory observations, or choose a sequence of actions, or collaborate with a number of agents, the data-acquisition task often involves inherent combinatorial structures and is fundamentally discrete. Even though discrete optimization problems are generally hard, prior work has shown that many data-acquisition problems admit a key structural property called submodularity. Due to recent breakthroughs in exploiting submodularity for discrete optimization, we now have efficient and provable algorithms for special cases of the data-acquisition problem. However, designing discrete and submodular optimization methodologies that are capable of adapting to dynamic and uncertain environments requires a quantum leap in the following three main directions, as aimed by the PI through this program: (i) developing foundational tools for discrete optimization in complex environments addressing uncertainty, resiliency, and unknown dynamics; (ii) designing polices that sequentially select the most informative data and observations while learning and adapting to the environment; (iii) developing cooperative strategies among multiple agents to jointly achieve similar goals as in (i),(ii).

15. SUBJECT TERMS

16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES
a. REPORT	b. ABSTRACT	c. THIS PAGE		
U	U	U	UU	9

19a. NAME OF RESPONSIBLE PERSON	19b. PHONE NUMBER *(Include area code)*
WARREN ADAMS	00000000

Standard Form 298 (Rev.5/2020)
Prescribed by ANSI Std. Z39.18

Data Acquisition in Dynamic Environments: A Submodular Perspective

Final Report
Young Investigator Research Program
Topic Area: Optimization and Discrete Mathematics
Program Manager: Dr. Warren Adams

PI: Hamed Hassani
Department of Electrical and Systems Engineering
University of Pennsylvania

Overview

While the recent advances in Artificial Intelligence have mainly relied on the availability of a wealth of centralized data, a fundamental challenge in many DoD-relevant applications is to acquire high-quality data at minimal cost. These applications range from robotic sensing and autonomous planning to experimental design and active learning; furthermore, such applications often take place in unknown or even adversarial environments, in which data is highly limited and precious. More specifically, each observation may significantly impact our ability to learn and operate in unknown and dynamic environments. Moreover, dealing with complex real-world environments requires a paradigm shift from the existing static, model-aware data acquisition approaches to methods that learn adaptively and are robust against imperfect, stochastic and evolving knowledge.

Whether we select a bunch of sensory observations, or choose a sequence of actions, or collaborate with a number of agents, the data-acquisition task often involves inherent combinatorial structures and is fundamentally discrete. Even though discrete optimization problems are generally hard, prior work has shown that many data-acquisition problems admit a key structural property called submodularity. Due to recent breakthroughs in exploiting submodularity for discrete optimization, we now have efficient and provable algorithms for special cases of the data-acquisition problem. However, designing discrete and submodular optimization methodologies that are capable of adapting to dynamic and uncertain environments requires a quantum leap in the following three main directions, as aimed by the PI through this program: (i) developing foundational tools for discrete optimization in complex environments addressing uncertainty, resiliency, and unknown dynamics; (ii) designing polices that sequentially select the most informative data and observations while learning and adapting to the environment; (iii) developing cooperative strategies among multiple agents to jointly achieve similar goals as in (i),(ii).

In this report we will describe our efforts within the three years duration towards the goals mentioned above. In particular, we will illustrate the following:

(i) Introducing and developing a new method to efficiently solve submodular optimization problems through learning from prior experience and data and adapting to new environments (the framework is called submodular meta-learning); The results have been presented at NeurIPS 2021 [1].

(ii) Devising new algorithms and cooperative strategies for distributed submodular problems; The results have been presented in the control/dynamic/learning community (L4DC 2021, [2]).

(iii) Introducing the class of convex-submodular minimax problems, where the problem-objective is convex with respect to the continuous variable and submodular with respect to the discrete variable, and developed a principled study of the problem from both theoretical and algorithmic perspectives; The results have been presented as an oral presentation at AISTATS 2022 [3]

(iv) designing efficient methodologies for multi-agent minmax learning problems in the presence of worst-case Byzantine adversarial agents; The results of this project were presented to the control/dynamic/learning community [4] (CDC 2022 as an invited paper).

(v) achieving fast optimization algorithms in the context of federated/distributed learning; The results appeared at NeurIPS 2021 [5].

(vi) a principled study (fundamental limits, efficient algorithms) for multi-agent sequential decision making problems in the presence of worst-case Byzantine adversarial agents. The results appeared at NeurIPS 2022 [6].

(vii) Studying effects of delay and asynchronous updates within stochastic approximation schemes, investigating non-asymptotic convergence rates under Markovian noise; The results will be presented at AISTATS 2024 [7].

In what follows, we will describe in more detail the scientific outcomes of this project.

Submodular Meta-Learning. Many applications in artificial intelligence necessitate exploiting prior data and experience to enhance quality and efficiency on new tasks. This is often manifested through a set of tasks given in the training phase from which we can learn a model or representation that can be used for new unseen tasks in the test phase. In this regard, Meta-learning aims at exploiting the data from the available tasks to learn model parameters or representation that can be later used to perform well on new unseen tasks, in particular, when we have access to limited data and computational power at the test time. By now, there are several formulations for Meta-learning, but perhaps one of the most successful ones is the Model-Agnostic Meta-Learning (MAML). In MAML, we aim to train the model parameters such that applying a few steps of gradient-based updates with a small number of samples from a new task would perform well on that task. MAML can also be viewed as a way to provide a proper initialization, from which performance on a new task can be optimized after a few gradient-based updates. Alas, this scheme only applies to settings in which the decision variable belongs to a continuous domain and can be adjusted using gradient-based methods at the test time.

Our goal in [1] is to extend the methodology of MAML to the discrete setting. We consider a setting that our decision variable is a discrete set, and our goal is to come up with a good initial set that can be quickly adjusted to perform well over a wide range of new tasks. In particular, we focus on submodular maximization to represent the tasks which is an essential class of discrete optimization. There are numerous applications where the submodular meta-learning framework can be applied to find a personalized solution for each task while significantly reducing the computation load. In general, most recommendation tasks can be cast as an instance of this setting.

Problem Formulation. We consider a family of tasks $\mathcal{T} = \{\mathcal{T}_i\}_{i\in\mathcal{I}}$, where the set $\mathcal{I}$ could be of infinite size. Each task $\mathcal{T}_i$ is represented via a set function $f_i : 2^V \to \mathbb{R}_+$ that measures the reward of a set $S \subseteq V$ for the i-th task, and performing the task $\mathcal{T}_i$ would mean to maximize the function f_i subject to a given constraint. For instance, in a recommender system where we aim to recommend a subset of the items to the users, the set $\mathcal{I}$ denotes the set of all the possible users and selecting which items to recommend to a user $i \in \mathcal{I}$ is viewed as the task $\mathcal{T}_i$. Moreover, the function f_i encodes the users satisfaction, i.e., $f_i(S)$ quantifies how suitable the set of items S is for user i. Taking a statistical perspective, we assume that the tasks $\mathcal{T}_i$ occur according to a possibly unknown probability distribution $i \sim p$. In [1], we focus on the case where the functions f_i are monotone and submodular set functions and each task $\mathcal{T}_i$ amounts to maximizing f_i under the k-cardinality constraint.

We assume access to a collection of *training* tasks $\{\mathcal{T}_i\}_{i=1}^m$. These are the tasks that we have already experienced, i.e., they correspond to the users that we have already seen. Formally, this means that for each training task $\mathcal{T}_i$, we assume knowledge of the corresponding function f_i. In our formulation, each of the training tasks is assumed to be generated i.i.d. according to the distribution p. Indeed, eventually we aim to optimize performance at *test* time, i.e., obtain the best performance for new and unseen tasks generated independently from the distribution p. For instance, in our recommendation setting, test tasks correspond to new users that will arrive in the future. Our goal is to use the training tasks to reduce the computation load at test time.

As we discussed so far, when computational power is limited at test time, it makes sense to divide the process of choosing the best decision between training and test phases. To be more specific, in the training phase, we choose a subset of elements from the ground set that would perform over the training tasks, and then select (or optimize) the remaining elements at the test time *specifically* with respect to the task at hand. To state this problem, consider $S_{tr} \subseteq V$ with cardinality $|S_{tr}| = l$, where $l < k$, as the initial set that we aim to find at the training phase, and the set S_i that we add to the initial set S_{tr} at test time. Hence, the problem of interest can be written as

$$\max_{S_{tr} \in V, |S_{tr}| \leq l} \mathbb{E}_{i \sim p} \Big[\max_{S_i \in V, |S_i| \leq k-l} f_i(S_{tr} \cup S_i) \Big], \tag{1}$$

Note that the critical decision variable that we need to find is S_{tr} which is the best initial subset of size l overall all possible choices of task when a best subset of size $k - l$ is added to that. In fact, if we define $f_i'(S_{tr}) := \max_{S_i \in V, |S_i| \leq k-l} f_i(S_{tr} \cup S_i)$, then we can rewrite the problem in (1) as

$$\max_{S_{tr} \in V, |S_{tr}| \leq l} \mathbb{E}_{i \sim p} \ \left[f_i'(S_{tr}) \right]. \tag{2}$$

As described previously, we often do not have access to the underlying probability distribution p of the tasks, and we instead have access to a large number of sampled tasked that are drawn independently according to p. Hence, instead of solving (1), we solve its sample average approximation given by

$$\max_{S_{tr} \in V, |S_{tr}| \leq l} \ \frac{1}{m} \sum_{i=1}^{m} \ \left[\max_{S_i \in V, |S_i| \leq k-l} f_i(S_{tr} \cup S_i) \right] \ = \ \max_{S_{tr} \in V, |S_{tr}| \leq l} \ \frac{1}{m} \sum_{i=1}^{m} \ \left[f_i'(S_{tr}) \right], \tag{3}$$

where m is the number of tasks in the training set which are sampled according to p.

Results. In [1], we present computationally efficient deterministic and randomized meta-greedy algorithms for solving Problem (3) with provable guarantees. When the tasks are monotone and submodular, we prove that the solution obtained by the deterministic algorithm is at least 0.53-optimal, and the solution of the randomized algorithm is $(1 - 1/e - o(1))$-optimal in expectation, where the $o(1)$ term vanishes by the size of the solution. These guarantees are obtained by introducing new techniques, despite that the meta-learning objective is *not* submodular.

We further studied the performance of our proposed meta-learning framework and algorithms for movie recommendation and ride-sharing problems. Our experiments illustrate that the solution of our proposed meta-learning scheme, which chooses a large portion of the solution in the training phase and a small portion adaptively at test time, is very close to the solution obtained by choosing the entire solution at the test time when a new task is revealed.

Submodular Maximization with Distributed Constraints. Recently, the need has arisen to design algorithms that distribute decision making among a collection of agents or computing devices. This need has been motivated by problems from statistics, machine learning and robotics. Among such problems, the problem of data/information acquisition is highly relevant to our research program. Inherent to such problems is an underlying optimization problem that can be expressed as

$$\begin{aligned} &\text{maximize } f(S) \\ &\text{s.t.} \quad S \subseteq \mathcal{Y} \text{ and } S \in \mathcal{I} \end{aligned} \tag{4}$$

where f is a submodular set function (i.e. it has a diminishing-returns property), $\mathcal{Y}$ is a finite set of all decision variables, and $\mathcal{I}$ is a family of allowable subsets of $\mathcal{Y}$. In words, the goal of (4) is to pick a set S from the family of allowable subsets $\mathcal{I}$ that maximizes the submodular set function f. A wide class of relevant objective functions such as mutual information and weighted coverage are submodular; this has motivated a growing body of work surrounding submodular optimization problems.

Intuitively, it is useful to think of the problem in (4) as a distributed n-player game. In this game, each player or agent has a distinct local strategy set of actions. The goal of the game is for each agent to choose at most one action from its own strategy set to maximize a problem-specific notion of reward. Therefore, the problem is *distributed* in the sense that agents can only form a control policy with the actions from their local, distinct strategy sets. To maximize reward, agents are allowed to communicate with their direct neighbors in a bidirectional communication graph. In this way, we might think of these agents as robots that collectively aim to solve a coverage problem in an unknown environment by communicating their sensing actions to their nearest neighbors. Throughout this work, we will refer to this multi-agent game example to elucidate our results.

In [2], our aim is to study problem (1) in a *distributed* setting, which we will formally introduce in the following; this setting differs considerably from the *centralized* setting, which has been studied thoroughly in past work. Notably, the distributed setting admits a more challenging problem because agents can only communicate locally with respect to a communication graph. Therefore designing an efficient communication scheme among agents is a concomitant requirement for the distributed setting, whereas in the centralized setting, there is no such desideratum.

Problem Formulation. The specific problem that we considered in [2] is submodular maximization subject to a distributed partition matroid constraint. Consider a collection of n agents that form the set $\mathcal{A} = \{1, \dots, n\}$. Let $f : 2^{\mathcal{Y}} \mapsto \mathbb{R}_+$ be a normalized and monotone submodular set function and let $\mathcal{Y}_1, \dots, \mathcal{Y}_n$ be a pairwise disjoint partition of a finite ground set $\mathcal{Y}$, wherein each agent $i \in \mathcal{A}$ can only choose actions from its local strategy set $\mathcal{Y}_i$. Furthermore, consider the partition matroid $(\mathcal{Y}, \mathcal{I})$, where

$$\mathcal{I} := \{S \subseteq \mathcal{Y} \ : \ |\mathcal{Y}_i \cap S| \leq 1 \text{ for } i = 1, \dots, n\}. \tag{5}$$

The problem of submodular maximization subject to a distributed partition matroid constraint is to maximize f by selecting a set $S \subseteq \mathcal{Y}$ from the family of allowable subsets so that $S \in \mathcal{I}$. In effect, the distributed partition matroid constraint enforces that each agent $i \in \mathcal{A}$ can choose at most one action from its local strategy set $\mathcal{Y}_i$. Note that in this setting, each agent can only choose actions from its own local strategy set. Therefore, this problem is distributed in the sense that agents can only determine the actions taken by other agents by directly communicating with one another.

The agents aim to solve Problem (4) subject to the distributed constrain (5) by collaborating/communication with each other. The inter-agent communication structure is given as follows. The agents $i \in \mathcal{A} = \{1, \dots, n\}$ share their decision variables with a small subset of *local* agents in $\mathcal{A}$. To encode the notion of locality, suppose that each agent $i \in \mathcal{A}$ is a node in a bidirectional *communication graph* $\mathcal{G} = (\mathcal{A}, \mathcal{E})$ in which $\mathcal{E}$ denotes the set of edges. Given this structure, we assume that each agent $i \in \mathcal{A}$ can only communicate its decision variable with its direct neighbors in $\mathcal{G}$. The goal is to solve the problem using the minimum amount of communication among the agents.

Results. We developed in [2] a decentralized algorithm for solving Problem (4). At a high level, this algorithm involves updating each agent's local decision variable based on the aggregated belief of a small group of other agents about the best control policy. In essence, inter-agent communication within small groups of agents facilitates local decision making. Our proposed algorithm is based on the a continuous extension of problem (4) with constraint (5) which is distributed across the agents. Using this equivalent extension, we can solve the problem in the continuous domain, through an iterative message passing procedure across the agents over the graph $\mathcal{G}$. The final solution is then rounded to a discrete solution satisfying the constraint in (5). We also offer an analysis of the proposed algorithm and prove that it achieves the tight $(1 - 1/e)$ approximation and that its error term vanishes at a linear rate.

Minimax Optimization: The Case of Convex-Submodular. Minimax optimization has been central in addressing various applications in machine learning, game theory, and control theory. Prior literature has thus far mainly focused on studying such problems in the continuous domain, e.g., convex-concave minimax optimization is now understood to a significant extent. Nevertheless, minimax problems extend far beyond the continuous domain to mixed continuous-discrete domains or even fully discrete domains. In this paper, we study mixed continuous-discrete minimax problems where the minimization is over a continuous variable belonging to Euclidean space and the maximization is over subsets of a given ground set.

Our goal in [3] is to introduce and provide a principled study of the class of We introduce the class of convex-submodular minimax problems, where the objective is convex with respect to the continuous variable and submodular with respect to the discrete variable. Even though such problems appear frequently in machine learning applications, little is known about how to address them from algorithmic and theoretical perspectives. In summary, our results are as follows: For such problems, we first show that obtaining saddle points are hard up to any approximation, and thus introduce new notions of (near-) optimality. We then provide several algorithmic procedures for solving convex and monotone-submodular minimax problems and characterize their convergence rates, computational complexity, and quality of the final solution according to our notions of optimally. Our proposed algorithms are iterative and combine tools from both discrete and continuous optimization. Finally, we have shown how our methods can be applied to achieve significantly better performance in real-word applications such as designing adversarial attacks for item recommendation.

Problem Formulation. We introduce the following structured non convex-concave minimax problems, where the minimization variable is from a continuous domain and the maximization variable belongs to a discrete domain. Concretely, for a non-negative function $f : \mathbb{R}^d \times 2^V \to \mathbb{R}_+$, consider the minimax problem

$$\text{OPT} \triangleq \min_{\mathbf{x}\in\mathcal{X}} \max_{S\in\mathcal{I}} f(\mathbf{x}, S), \tag{6}$$

where $\mathbf{x}$ belongs to a convex set $\mathcal{X} \subset \mathbb{R}^d$ and S is a subset of the ground set V with n elements that is constrained to be inside a matroid $\mathcal{I}$. Given a fixed S, the function $f(\cdot, S)$ is convex with respect to the continuous (minimization) variable. Further, given a fixed $\mathbf{x}$, the function $f(\mathbf{x}, \cdot)$ is submodular with respect to the discrete (maximization) variable. We refer to this problem as *convex-submodular minimax problem*.

The convex-submodular minimax problem in (6) encompasses various applications. In particular, when convex models have to be learned while data points are selected or changed according to notions of summarization, diversity, and deletion. Examples include learning under data deletion, robust text classification, minimax curriculum learning, minimax supervised learning, and minimax active learning.

Results. In [3], we provide a principled study of the problem defined in (6), from both theoretical and algorithmic perspectives, when f is convex in the minimization variable and submodular as well as *monotone* in the maximization variable[1]. We introduce efficient iterative algorithms for solving this problem and develop a theoretical framework for analyzing such algorithms with guarantees on the quality of the resulting solutions according to the notions of optimality that we define. Next, we provide a detailed explanation of our results.

(i) Notions of (near-)optimality and hardness results. For minimax problems, the strongest notion of optimality is defined through saddle points or their approximate versions. We first provide a negative result that shows finding a saddle point or any approximate version of it (which we term as an (α, ϵ)-saddle point) is NP-hard for general convex-submodular problems. We thus introduce a slightly weaker notion of optimality that we call (α, ϵ)-approximate minimax solutions for Problem (6). Roughly speaking, the quality of the minimax objective at such solutions is at most $\frac{1}{\alpha}(\text{OPT} + \epsilon)$, and hence they are near-optimal when $\alpha < 1$. We then show that obtaining such solutions for $\alpha > 1 - 1/e$ is NP-hard. This is a non-trivial result that does not readily follow from known hardness results in submodular maximization. Consequently, we focus on efficiently finding solutions in the regime of $\alpha \leq (1 - 1/e)$. We present several algorithms that achieve this goal and theoretically analyze their complexity and quality of their solution.

[1]For completeness, a function $g : 2^V \to \mathbb{R}$ is called submodular if for any two subsets $S, T \subseteq V$ we have: $g(S \cap T) + g(S \cup T) \leq g(S) + g(T)$. Moreover, g is called monotone if for any $S \subseteq T$ we have $g(S) \leq g(T)$.

Table 1: Algorithms performance guarantee. Here c_f is the cost of single computation of f, $c_{P_{\mathbf{x}}}$ and c_P are cost of projection in $\mathcal{X}$ and $\mathcal{Y}$, $c_{\nabla_{\mathbf{x}} f}$ is the cost of computing gradient of f with respect to $\mathbf{x}$, and $c_{\nabla_{\mathbf{x}} F}$ and $c_{\nabla F}$ are the cost of computing gradient of multilinear extension F with respect to $\mathbf{x}$ and , respectively. k is the cardinality constraint ($|S| \leq k$) and n is size of the ground set $|V| = n$.

Alg.	Number of iterations	Approx. ratio	Cost per iteration	Card. const.	Matroid const.	Unbounded grad.
GG	$\mathcal{O}(1/\epsilon^2)$	$1-1/e$	$nk.c_f + c_{P_{\mathbf{x}}} + c_{\nabla_{\mathbf{x}} f}$	✓	✗	✗
GG	$\mathcal{O}(1/\epsilon^2)$	$1/2$	$nk.c_f + c_{P_{\mathbf{x}}} + c_{\nabla_{\mathbf{x}} f}$	✗	✓	✗
GRG	$\mathcal{O}(1/\epsilon^2)$	$1/2$	$(n+k)c_f + c_{\nabla_{\mathbf{x}} f + c_{P_{\mathbf{x}}}}$	✓	✗	✗
EGG	$\mathcal{O}(1/\epsilon^2)$	$1-1/e$	$2nk.c_f + 2c_{P_{\mathbf{x}}} + 2c_{\nabla_{\mathbf{x}} f}$	✓	✗	✓
EGG	$\mathcal{O}(1/\epsilon^2)$	$1/2$	$2nk.c_f + 2c_{P_{\mathbf{x}}} + 2c_{\nabla_{\mathbf{x}} f}$	✓	✓	✓
EGRG	$\mathcal{O}(1/\epsilon^2)$	$1/2$	$2(n+k)c_f + 2c_{P_{\mathbf{x}}} + 2c_{\nabla_{\mathbf{x}} f}$	✓	✗	✗
EGCE	$\mathcal{O}(1/\epsilon)$	$1/2$	$2c_{P_{\mathbf{x}}} + 2c_P + 2c_{\nabla_{\mathbf{x}} F} + 2c_{\nabla F}$	✓	✓	✓

(ii) Algorithms with guarantees on convergence rate, complexity, and solution quality. Our proposed algorithms are as follows (see also Table 1): (i) *Greedy-based methods.* We first present Gradient-Greedy (GG), a method alternating between gradient descent for minimization and greedy for maximization. We further introduce Extra-Gradient-Greedy (EGG) that uses an extra-gradient step instead of gradient step for the minimization variable. We prove that both algorithms achieve a $((1-1/e), \epsilon)$-approximate minimax solution after $\mathcal{O}(1/\epsilon^2)$ iterations when $\mathcal{I}$ is a cardinality constraint. Importantly, EGG does not require the bounded gradient norm condition as opposed to GG. (ii) *Replacement greedy-based methods.* The greedy-based methods require $\mathcal{O}(nk)$ function computations at each iteration. To improve this complexity, we present alternating methods that use replacement greedy for the maximization part to reduce the cost of each iteration to $\mathcal{O}(n)$. The Gradient Replacement-Greedy (GRG) algorithm achieves a $(1/2, \epsilon)$-approximate minimax solution after $\mathcal{O}(1/\epsilon^2)$ iterations and Extra-Gradient Replacement-Greedy (EGRG) achieves a $(1/2, \epsilon)$-approximate minimax solution after $\mathcal{O}(1/\epsilon^2)$, when $\mathcal{I}$ is a cardinality constraint. (iii) *Continuous extension-based methods.* Note that all mentioned methods achieve a convergence rate of $\mathcal{O}(1/\epsilon^2)$. To improve this convergence rate, we further introduce the extra-gradient on continuous extension (EGCE) method that runs extra-gradient update on the continuous extension of the submodular function. We show that EGCE is able to achieve an $(1/2, \epsilon)$-approximate minimax solution after at most $\mathcal{O}(1/\epsilon)$ iterations, when $\mathcal{I}$ is a general matroid constraint.

Distributed Statistical Min-Max Learning in the Presence of Byzantine Agents. Recent years have witnessed a growing interest in the topic of min-max optimization, owing to its relevance in the context of generative adversarial networks (GANs), robust control and optimization, and reinforcement learning. Motivated by this line of work, we consider a multi-agent minmax learning problem, and focus on the emerging challenge of contending with worst-case Byzantine adversarial agents in such a setup. By drawing on recent results from robust statistics, we design a robust distributed variant of the extragradient algorithm - a popular algorithmic approach for minmax optimization. Our main contribution is to provide a crisp analysis of the proposed robust extra-gradient algorithm for smooth convex-concave and smooth strongly convex-strongly concave functions.

In summary, in [4] we establish statistical rates of convergence to approximate saddle points. Our rates are near-optimal, and reveal both the effect of adversarial corruption and the benefit of collaboration among the non-faulty agents. Notably, [4] is the first work to provide formal theoretical guarantees for large-scale distributed min-max learning in the presence of adversarial agents.

Problem Formulation. We consider a min-max learning problem of the form

$$\min_{x \in \mathcal{X}} \max_{y \in \mathcal{Y}} f(x, y) \triangleq \mathbb{E}_{\xi \sim \mathcal{D}}[F(x, y; \xi)]. \tag{7}$$

Here, $\mathcal{X}$ and $\mathcal{Y}$ are convex, compact sets in $\mathbb{R}^n$ and $\mathbb{R}^m$, respectively; $x \in \mathcal{X}$ and $y \in \mathcal{Y}$ are model parameters; ξ is a random variable representing a data point sampled from the distribution $\mathcal{D}$; and $f(x, y)$ is the population function corresponding to the stochastic function $F(x, y; \xi)$. Throughout this paper, we assume that $f(x, y)$ is continuously differentiable in x and y, and is *convex-concave* over $\mathcal{X} \times \mathcal{Y}$. Specifically, $f(\cdot, y) : \mathcal{X} \to \mathbb{R}$ is convex for every $y \in \mathcal{Y}$, and $f(x, \cdot) : \mathcal{Y} \to \mathbb{R}$ is concave for every $x \in \mathcal{X}$. Our goal is to find a saddle point (x^*, y^*) of $f(x, y)$ over the set $\mathcal{X} \times \mathcal{Y}$, where a saddle point is defined as a vector pair $(x^*, y^*) \in \mathcal{X} \times \mathcal{Y}$ that satisfies

$$f(x^*, y) \leq f(x^*, y^*) \leq f(x, y^*), \forall x \in \mathcal{X}, y \in \mathcal{Y}. \tag{8}$$

The min-max optimization problem described above features in a variety of applications: from classical developments in game theory and online learning, to robust optimization and reinforcement learning. More recently, in the context of machine learning, min-max problems have found important applications in training generative adversarial networks (GANs), and in robustifying deep neural networks against adversarial attacks. Motivated by this recent line of work, we consider a min-max learning problem of the form in Eq. (7), where the data samples required for finding a saddle-point are distributed across multiple devices (agents). Specifically, we focus on a large-scale distributed setup comprising of M agents, each of which can access i.i.d. data samples from the distribution $\mathcal{D}$. The agents collaborate under the orchestration of a central server to compute an approximate saddle point of statistical accuracy higher relative to the setting when they act alone. The intuition here is simple: since all agents receive data samples from the *same* distribution, exchanging information via the server can help reduce the randomness (variance) associated with these samples. An example of the above setup that aligns with the modern federated learning paradigm is one where multiple devices (e.g., cell phones or tablets) collaborate via a server to train a robust statistical model.

To reap the benefits of collaboration in modern distributed computing systems, one needs to contend with the critical challenge of *security*. In particular, this challenge arises from the fact that the individual agents in such systems are easily susceptible to adversarial attacks. In fact, unless appropriately accounted for, even a single malicious agent can severely degrade the overall performance of the system by sending corrupted messages to the central server.

Objective. Thus, given the emerging need for security in large-scale computing, *our objective in [4] is to design an algorithm that achieves near-optimal statistical performance in the context of distributed min-max learning, while being robust to worst-case attacks*. To that end, we consider a setting where a fraction of the agents is Byzantine. Each Byzantine agent is assumed to have complete knowledge of the system and learning algorithms; moreover, leveraging such knowledge, the Byzantine agents can send arbitrary messages to the server and collude with each other.

Challenges. Even in the absence of noise or attacks, it is known that algorithms such as gradient descent ascent (`GDA`) can diverge for simple convex-concave functions. We have to contend with both noise (due to our statistical setup) *and* worst-case attacks - this makes the analysis for our setting non-trivial. In particular, the adversarial agents can introduce complex probabilistic dependencies across iterations that need to be carefully accounted for; we do so in this work by making the following contributions.

Results. In summary, the results of [4] are:

• *Problem.* Given the importance and relevance of security, several recent works have studied distributed optimization/learning in the face of adversarial agents. However, we are unaware of any analogous paper for adverarially-robust distributed *min-max* learning. Our work closes this gap.

• *Algorithm.* We develop an algorithm for finding an approximate saddle point to the min-max learning problem in Eq. (7), subject to the presence of Byzantine agents. Our proposed algorithm - called Robust Distributed Extra-Gradient (`RDEG`) - brings together two separate algorithmic ideas: (i) the classical extra-gradient algorithm due to Korpelevich that has gained a lot of popularity due to its empirical performance in training GANs, and (ii) the recently proposed univariate trimmed mean estimator due to Lugosi and Mendelson.

• *Theoretical Results.* Our main contribution is to provide a rigorous theoretical analysis of the performance of `RDEG` for smooth convex-concave and smooth strongly convex-strongly concave settings. In each case, we

establish that as long as the fraction of corrupted agents is "small", RDEG guarantees convergence to approximate saddle points at *near-optimal* statistical rates with high probability. The rates that we derive precisely highlight the benefit of collaboration in effectively reducing the variance of the noise model. At the same time, they indicate the (unavoidable) additive bias introduced by adversarial corruption. Finally, an immediate benefit of such an analysis is that one can build on it for the more challenging nonconvex-nonconcave setting as future work.

Fast Optimization Methodologies for Federated Learning. In [5], we consider a standard federated learning (FL) setup where a group of clients periodically coordinate with a central server to train a statistical model. Our main result is to develop a general algorithmic framework called FedLin to tackle some of the key challenges intrinsic to FL, namely objective heterogeneity, systems heterogeneity, and infrequent and imprecise communication. Our framework is motivated by the observation that under these challenges, various existing FL algorithms suffer from a fundamental speed-accuracy conflict: they either guarantee linear convergence but to an incorrect point, or convergence to the global minimum but at a sub-linear rate, ie, fast convergence comes at the expense of accuracy. In contrast, when the clients' local loss functions are smooth and strongly convex, we show that FedLin guarantees linear convergence to the global minimum, despite arbitrary objective and systems heterogeneity. We then establish matching upper and lower bounds on the convergence rate of FedLin that highlight the effects of infrequent, periodic communication. Finally, we show that FedLin preserves linear convergence rates under aggressive gradient sparsification, and quantify the effect of the compression level on the convergence rate.

It is worth noting that FedLin is the first to provide tight linear convergence rate guarantees, and constitutes the first comprehensive analysis of gradient sparsification in FL.

Collaborative Learning with Adversarial Agents. Given our focus on complex environments, we began as the first step to study the simplest (yet foundational) instance of sequential decision making problems, namely Bandits, in complex environments [6]. In detail, we consider a linear stochastic bandit problem involving M agents that can collaborate via a central server to minimize regret. A fraction a of these agents are adversarial and can act arbitrarily, leading to the following tension: while collaboration can potentially reduce regret, it can also disrupt the process of learning due to adversaries. We have succeeded in providing a fundamental understanding of this tension by designing new algorithms that balance the exploration-exploitation trade-off via carefully constructed robust confidence intervals. We also complement our algorithms with tight analyses. First, we develop a robust collaborative phased elimination algorithm that achieves a regret of order $(a+1/\sqrt{m})\sqrt{dT}$ for each good agent; here, d is the model-dimension and T is the horizon. For small a, our result thus reveals a clear benefit of collaboration despite adversaries. Moreover, using an information-theoretic argument, we have been able to prove a matching lower bound, thereby providing the first set of tight, near-optimal regret bounds for collaborative linear bandits with adversaries. Furthermore, by leveraging recent advances in high-dimensional robust statistics, we significantly extend our algorithmic ideas and results to (i) the generalized linear bandit model that allows for non-linear observation maps; and (ii) the contextual bandit setting that allows for time-varying feature vectors.

Student and Postdoc Training

This project led to successful training of two PhD students, Arman Adibi (who is now a postdoc at Princeton University) and Donghwan Lee (who successfully defended his thesis and will work in the financial sector), as well as one postdoctoral researcher, Aritra Mitra, who is now an assistant professor at the North Carolina State University.

Final Remarks

In summary, we are happy to report that this project made significant progress towards all its three main goals listed at the beginning of this report. Moreover, this project led to two successful PhD theses and trained a postdoctoral researcher. Among the three trainees of this project, two have decided to pursue an academic path. We are very thankful to the AirForce Office of Scientific Research for supporting this project and for giving us the opportunity to achieve the above ambitious goals.

References

[1] A. Adibi, A. Mokhtari, and H. Hassani, "Submodular meta-learning," *Advances in Neural Information Processing (NeurIPS)*, pp. 150–162, 2020.

[2] A. Robey, A. Adibi, B. Schlotfeldt, G. J. Pappas, and H. Hassani, "Optimal algorithms for submodular maximization with distributed constraints," in *Learning for Dynamics and Control*, pp. 150–162, PMLR, 2021.

[3] A. Adibi, A. Mokhtari, and H. Hassani, "Minimax optimization: The case of convex-submodular," *Artificial Intelligence and Statistics Conference (AISTATS)*, 2022.

[4] A. Adibi, A. Mitra, G. J. Pappas, and H. Hassani, "Distributed statistical min-max learning in the presence of byzantine agents," *arXiv preprint arXiv:2204.03187*, 2022.

[5] A. Mitra, R. Jaafar, G. J. Pappas, and H. Hassani, "Linear convergence in federated learning: Tackling client heterogeneity and sparse gradients," *Advances in Neural Information Processing Systems*, vol. 34, pp. 14606–14619, 2021.

[6] A. Mitra, A. Adibi, G. J. Pappas, and H. Hassani, "Collaborative linear bandits with adversarial agents: Near-optimal regret bounds," *arXiv preprint arXiv:2206.02834*, 2022.

[7] A. Adibi, N. Dal Fabbro, L. Schenato, S. Kulkarni, H. V. Poor, G. J. Pappas, H. Hassani, and A. Mitra, "Stochastic approximation with delayed updates: Finite-time rates under markovian sampling," in *International Conference on Artificial Intelligence and Statistics*, pp. 2746–2754, PMLR, 2024.