AFRL-AFOSR-VA-TR-2021-0100

Efficient Mathematical Methods for the Optimization of Large and Complex Systems

Javad Lavaei
REGENTS OF THE UNIVERSITY OF CALIFORNIA
2150 SHATTUCK AVE RM 313
BERKELEY, CA,
US

08/17/2021
Final Technical Report

DISTRIBUTION A: Distribution approved for public release.

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

REPORT DOCUMENTATION PAGE	*Form Approved* *OMB No. 0704-0188*

1. REPORT DATE *(DD-MM-YYYY)*	**2. REPORT TYPE**	**3. DATES COVERED** *(From - To)*
17-08-2021	Final	15 Feb 2017 - 31 Dec 2020

4. TITLE AND SUBTITLE
Efficient Mathematical Methods for the Optimization of Large and Complex Systems

5a. CONTRACT NUMBER

5b. GRANT NUMBER
FA9550-17-1-0163

5c. PROGRAM ELEMENT NUMBER
61102F

6. AUTHOR(S)
Javad Lavaei

5d. PROJECT NUMBER

5e. TASK NUMBER

5f. WORK UNIT NUMBER

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)
REGENTS OF THE UNIVERSITY OF CALIFORNIA
2150 SHATTUCK AVE RM 313
BERKELEY, CA
US

8. PERFORMING ORGANIZATION REPORT NUMBER

9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)
AF 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-2021-0100

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

13. SUPPLEMENTARY NOTES

14. ABSTRACT
This project studies complex large-scale optimization problems, numerical algorithm, and applications in machine learning, optimal control, transportation, and power systems. A summary of some of our results developed during the last performance period is provided below:
Projection-free online convex optimization: We considered structured online convex optimization (OCO) with bandit feedback, where either the loss function is smooth or the constraint set is strongly convex. Projection- free methods are among the most popular and computationally efficient algorithms for solving this problem, mainly due to their ability to handle convex constraints appearing in machine learning for which computing projections is often impractical in high-dimensional settings. Despite the improved regret bound results for the full-information setting where the gradients of the functions are readily available, it remains unclear whether simple projection-free zero-order algorithms become more efficient for structured OCO problems in the case when multiple function values can be sampled at each time instance.

15. SUBJECT TERMS

16. SECURITY CLASSIFICATION OF:			**17. LIMITATION OF ABSTRACT**	**18. NUMBER OF PAGES**	**19a. NAME OF RESPONSIBLE PERSON**
a. REPORT	**b. ABSTRACT**	**c. THIS PAGE**			FARIBA FAHROO
U	U	U	UU	3	**19b. TELEPHONE NUMBER** *(Include area code)* 426-8429

This project studies complex large-scale optimization problems, numerical algorithm, and applications in machine learning, optimal control, transportation, and power systems. A summary of some of our results developed during the last performance period is provided below:

Projection-free online convex optimization: We considered structured online convex optimization (OCO) with bandit feedback, where either the loss function is smooth or the constraint set is strongly convex. Projection- free methods are among the most popular and computationally efficient algorithms for solving this problem, mainly due to their ability to handle convex constraints appearing in machine learning for which computing projections is often impractical in high-dimensional settings. Despite the improved regret bound results for the full-information setting where the gradients of the functions are readily available, it remains unclear whether simple projection-free zero-order algorithms become more efficient for structured OCO problems in the case when multiple function values can be sampled at each time instance. In this work, we developed some simple projection-free algorithms and proved that they indeed achieve the same improved regret bounds as the full-information case under various additional problem structures. This implies that leveraging the structural properties of the problem compensates for the lack of access to the gradients. Experiments on the online matrix completion revealed several attractive advantages of the proposed algorithms, including their simplicity, easy implementation, and effectiveness, as they outperform other competing algorithms.

Stochastic localization methods for discrete convex simulation optimization: We proposed a set of new algorithms based on stochastic localization methods for large-scale discrete simulation optimization problems with convexity structure. All proposed algorithms, with the general idea of "localizing" potential good solutions to an adaptively shrinking subset, are guaranteed with high probability to identify a solution that is close enough to the optimal given any precision level. Specifically, for one-dimensional large-scale problems, we proposed an enhanced adaptive algorithm with an expected simulation cost asymptotically independent of the problem scale, which is proved to attain the best achievable performance. For multi-dimensional large-scale problems, we proposed statistically guaranteed stochastic cutting-plane algorithms, the simulation costs of which have no dependence on model parameters such as the Lipschitz parameter, as well as low polynomial order of dependence on the problem scale and dimension.

Global and local analyses of nonlinear low-rank matrix recovery problems: The restricted isometry property (RIP) is a well-known condition that guarantees the absence of spurious local minima in low-rank matrix recovery problems with linear measurements. In this work, we introduced a novel property named bound difference property (BDP) to study low-rank matrix recovery problems with nonlinear measurements. Using RIP and BDP jointly, we first focused on the rank-1 matrix recovery problem, for which we proposed a new criterion to certify the nonexistence of spurious local minima over the entire space. We then analyzed the general case with an arbitrary rank and derived a condition to rule out the possibility of having a spurious solution in a ball around the true solution. The developed conditions lead to much stronger theoretical guarantees than the existing bounds on RIP.

Analysis of spurious local solutions of optimal control problems: One-shot optimization versus dynamic programming: Dynamic programming (DP) has a rich theoretical foundation and a broad range of applications, especially in the classic area of optimal control and the recent area of reinforcement learning (RL). Many optimal control problems can be solved as a single optimization problem, named one-shot optimization, or via a sequence of optimization problems using DP. However, the computation of their global optima often faces the NP-hardness issue due to the non-linearity of the dynamics and non-convexity of the cost, and thus only local optimal solutions may be obtained at best. Furthermore, in many cases arising in machine learning and model-free approaches, DP is the only viable choice, and therefore it is essential to understand when DP combined with a local search solver works. In this work, we introduced the notions of spurious local minimizers for the one-shot optimization and spurious local minimum policies for DP, and showed that there is a deep connection between them. In particular, we proved that under mild conditions the DP method using local search can successfully solve the optimal control problem to global optimality if and only if the one-shot optimization is free of spurious solutions. This result paves the way to understand the performance of local search methods in optimal control and RL.

Learning of dynamical systems under adversarial attacks: We studied the identification of a linear time-invariant dynamical systems affected by large-and-sparse disturbances modeling adversarial attacks or faults. Under the assumption that the states are measurable, we developed sufficient conditions for the recovery of the system matrices by solving a constrained lasso-type optimization problem. We showed how the input of the system should be designed to properly excite the system so that the constrained lasso problem correctly solves the system identification problem in an adversarial setting.

Adversarial attacks on computation of the policy iteration method: Adversarial attacks on Markov decision processes (MDPs) and reinforcement learning (RL) have been studied in the literature in the context of robust learning and adversarial game theory. In this work, we introduced a new notion of adversarial attacks on MDP and RL computation that is motivated by the emergence of edge computing. The large-scale computation of MDP and RL models in the form of value/policy iteration and Q-learning is being offloaded from agents to distributed servers, giving rise to edge reinforcement learning. By the inherently distributed nature of edge RL, the MDP/RL computation can be prone to adversarial attacks in different forms. We analyzed a probabilistic model of adversarial attacks on the computation of the modified policy iteration method in which the principal contraction property of the Bellman operator is undermined with a certain probability in iterations of the policy evaluation step of the aforementioned method. This can result in luring the agent to search among sub-optimal policies without improving the true values of policies. We proved that under certain conditions, the attacked modified policy iteration method can still converge to the vicinity of the optimal policy with high probability if the number of policy evaluation iterations is larger than a threshold that is logarithmic in the inverse of a desired precision. We also provided an upper bound on the number of iterations needed for the attacked modified policy iteration method to terminate, which holds with an associated confidence level.