Proceedings of Machine Learning Research vol 247:1–1, 2024 37th Annual Conference on Learning Theory

Mirror Descent Algorithms with Nearly Dimension-Independent Rates for Differentially-Private Stochastic Saddle-Point Problems

Tomás González TCGONZAL@ANDREW.CMU.EDU
Machine Learning Department, Carnegie Mellon University

Cristóbal Guzmán CRGUZMANP@MAT.UC.CL
Faculty of Mathematics and School of Engineering, Pontificia Universidad Católica de Chile
Google Research

Courtney Paquette COURTNEY.PAQUETTE@MCGILL.CA
Department of Mathematics and Statistics, McGill University
Google Deepmind

Editors: Shipra Agrawal and Aaron Roth

Abstract

We study the problem of differentially-private (DP) stochastic (convex-concave) saddle-points in the polyhedral setting. We propose (ε, δ)-DP algorithms based on stochastic mirror descent that attain nearly dimension-independent convergence rates for the expected duality gap, a type of guarantee that was known before only for bilinear objectives. For convex-concave and first-order-smooth stochastic objectives, our algorithms attain a rate of $\sqrt{\log(d)/n} + (\log(d)^{3/2}/[n\varepsilon])^{1/3}$, where d is the dimension of the problem and n the dataset size. Under an additional second-order-smoothness assumption, we improve the rate on the expected gap to $\sqrt{\log(d)/n} + (\log(d)^{3/2}/[n\varepsilon])^{2/5}$. Under this additional assumption, we also show, by using bias-reduced gradient estimators, that the duality gap is bounded by $\log(d)/\sqrt{n} + \log(d)/[n\varepsilon]^{1/2}$ with constant success probability. This result provides evidence of the near-optimality of the approach. Finally, we show that combining our methods with acceleration techniques from online learning leads to the first algorithm for DP Stochastic Convex Optimization in the polyhedral setting that is not based on Frank-Wolfe methods. For convex and first-order-smooth stochastic objectives, our algorithms attain an excess risk of $\sqrt{\log(d)/n} + \log(d)^{7/10}/[n\varepsilon]^{2/5}$, and when additionally assuming second-order-smoothness, we improve the rate to $\sqrt{\log(d)/n} + \log(d)/\sqrt{n\varepsilon}$. Instrumental to all of these results are various extensions of the classical Maurey Sparsification Lemma, which may be of independent interest.

Keywords: Differential Privacy, Stochastic Saddle Point Problem, Mirror Descent, Sparse Approximation

Acknowledgments

T. González's research was partially done as a Research Intern at Google DeepMind. C. Guzmán's research was partially supported by INRIA Associate Teams project, ANID FONDECYT 1210362 grant, ANID Anillo ACT210005 grant, and National Center for Artificial Intelligence CENIA FB210017, Basal ANID. C. Paquette is a Canadian Institute for Advanced Research (CIFAR) AI chair, Quebec AI Institute (MILA) and C. Paquette was supported by a Discovery Grant from the NSERC, NSERC CREATE grant INTER-MATH-AI, an external grant from Google, and FRQNT New University Researcher's Start-Up Program.

. Extended abstract. Full version appears as [arxiv reference,2403.02912]