CAREER: Robust Adaptive Optimization Algorithms for Differentially Private Learning

UNIVERSITY OF GEORGIA

Challenge:

- The use of differentially private optimization algorithms in practice is limited due to lack of understanding and theoretical tools to explain their performance.
- Practical private optimization algorithms need to provide reliable performance.

Solution:

- Developing a unified framework for systematic performance analysis
- Robustifying private algorithms by adding runtime adaptation of hyperparameters (e.g., privacy budget, step size, mini-batch size, etc.)
- Designing training algorithms for deep learning models, guided by theory

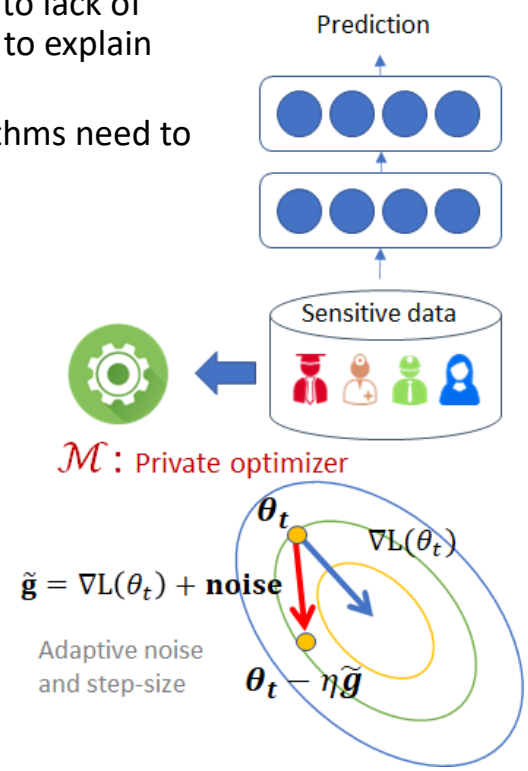

Scientific Impact:

- Providing a tool to understand the finite-time behavior of differentially private optimizers
- Formulating the algorithm performance as a function of important quantities: e.g., privacy budget, sample size, step-sizes, etc.
- Making private optimizers practical by improving their robustness to hyperparameter choice

Broader Impact and Broader Participation:

- Providing researchers in other domains a way to analyze sensitive data
- Companies with a large user base can use private optimizers to mitigate privacy concerns on applying cutting-edge ML techniques to user data
- Providing a research opportunity for undergrads and advanced trainings on the state-of-the art privacy-preserving techniques for graduate students

Award #1943046

PI: Jaewoo Lee (jwlee@cs.uga.edu)
Institution: University of Georgia