CAREER: Robust Adaptive Optimization Algorithms for Differentially Private Learning

PI: Jaewoo Lee, Department of Computer Science, University of Georgia

Objectives

1. To advance our understanding on the behavior of private optimization algorithms by providing a unified framework for performance analysis
2. To improve the performance reliability by increasing the robustness of private optimization algorithms
3. To extend the framework to deep learning models

Progress to date

- Studied the importance of runtime adaptivity of private optimization algorithms to the performance
  - ❑ Developed a stochastic line search method for choosing step size for differentially private algorithms [Chen and Lee 2020a]
  - ❑ adaptive ADMM algorithm [Chen and Lee 2020b]
- Extending ideas to deep learning models
  - ❑ Developed an alternative to BP algorithm: direct feedback alignment (DFA) algorithm [Lee and Kifer 2020]
  - ❑ Developed a fast gradient clipping method for training deep neural networks [Lee and Kifer 2021]
  - ❑ Differentially private normalizing flow models [Lee et al 2022]

Challenges

- Existing algorithms are often designed based on *heuristics* and restrictive assumptions, rather than being guided by *principled theory*.
- Asymptotic utility analysis is often provided to understand the performance of private optimization algorithms, but they are *not practical* as the analysis describes the behavior of optimizer when $n \to \infty$ and in expectation over all noise.
- Differentially private optimizers need to provide *reliable* performance. Their performance should *not* be sensitive to hyperparameter choices (e.g., privacy budget, number of iterations, step-sizes).

Scientific Impact

- The unified performance analysis framework will characterize the private algorithm's *convergence behavior* in response to noise added by the privacy mechanism.
- The project studies two important *practical aspects* of private optimizers: (i) finite-time performance and (ii) robustness to hyperparameter variation.
- The project will provide some *insight* on considerations for training deep neural networks under differential privacy.

Solutions

1. A unified framework for analyzing performance of differentially private optimizers and construct a Lyapunov function that demonstrate the stability of system.
   - Viewing an optimization algorithm as a linear time-invariant dynamical system

$$x_{t+1} = y_t - \eta \nabla f(y_t)$$
$$y_t = (1+\beta)x_t - \beta x_{t-1}$$

Example: Nesterov's accelerated gradient method

⇩

$$\underbrace{\begin{bmatrix} x_{t+1} - x^* \\ x_t - x^* \end{bmatrix}}_{\xi_{t+1}} = \left(\begin{bmatrix} 1+\beta & -\beta \\ 1 & 0 \end{bmatrix} \otimes I_p \right) \underbrace{\begin{bmatrix} x_t - x^* \\ x_{t-1} - x^* \end{bmatrix}}_{\xi_t} + \left(\begin{bmatrix} -\eta \\ 0 \end{bmatrix} \otimes I_p \right) \nabla f(y_t)$$

   - $\xi^T P \xi \to 0$ means $x_t \to x^*$.

2. Robustifying private optimization algorithms through algorithmic adaptivity
   - Instead of specifying per-iteration privacy budgets in advance, adaptively allocate the budget ϵ_t based on the usefulness of noisy gradient $\nabla f(x_t)$ during the runtime.
   - Add line-search functionality to private optimizers using the Sparse Vector Technique.
   - Dynamically adjust the amount of noise added to the gradient
3. Designing efficient private optimizers for deep learning models
   - Exploring alternatives to the gradient clipping: bounding the sensitivity through architectural constraints (e.g., orthogonality, scale invariance, and so on).
   - Exploring alternative to the backpropagation algorithm such as DFA

Broader Impact

- The performance analysis framework can enhance our understanding of private optimization algorithms behavior.
- The adaptive private optimizer developed in this project provides researchers in other disciplines and practitioners an easy way to apply differential privacy to their data analyses.
- The tools developed in this project can help our society at large by mitigating privacy concerns on machine learning applications.

Education

- Adding research components to undergraduate & graduate courses: the research outcome of project has been used as class project materials for the courses (CSCI 4260 & CSCI 8960) PI is teaching.
- The project provided research opportunities to undergraduate students: 3 undergrads participated in the PI's research program and the outcome was presented as a poster.
- Promoting computer science research to high school students.

Broadening participation

- A 3-day summer camp for local high school students will be held on July 13-15, 2022. The expected number of participants is around 25 – 30.
- The project recruited 3 undergraduate research assistants through UGA's CURO (Center for Undergraduate Research Opportunity) program.
- The project supported 3 CS Ph.D. students and the research outcome will serve as the basis for their dissertation.