Differentially Private Stochastic Coordinate Descent

Georgios Damaskinos,[1,2] Celestine Mendler-Dünner,[2,3]
Rachid Guerraoui,[1] Nikolaos Papandreou,[2] Thomas Parnell [2]

1 EPFL
2 IBM
3 UC Berkeley

AAAI 2021

Problem

SCD is **popular** in both Academia and Industry

- 154 research articles with "coordinate descent" in the title since 2019
- Default solver for:
  Scikit-Learn, TensorFlow, Liblinear, IBM Snap-ML

Why so popular ?

✓ Low tuning cost (no learning rate)
✓ Often favorable convergence guarantees
➢ In particular for GLMs

SCD applications involve **sensitive data**

- healthcare
- finance
- social media
- recommenders

...

*Can SCD maintain its **benefits** alongside **strong privacy** guarantees ?*

DP-SCD

Challenge

Differential privacy requires
independent noise added to **α** and **w**
=> No consistency: $\mathbf{w} \neq \mathbf{X}^T \cdot \boldsymbol{\alpha}$

1. Convergence guarantees ?
2. Competitive privacy-utility trade-off ?

Design

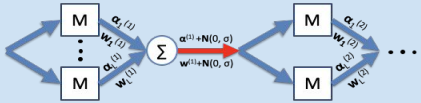

➔ Parallel updates (mini-batch)
➔ Update scaling

Convergence

Consistency *holds in expectation*

Method	Perturbation	Utility Bound
(Zhang et al. 2017)	Output	$\mathcal{O}\left(\frac{m}{n^2\epsilon^2}\right)$
(Chaudhuri and Monteleoni 2009) (Chaudhuri, Monteleoni, and Sarwate 2011)	Inner (objective)	$\mathcal{O}\left(\frac{m}{n^2\epsilon^2}\right)$
(Wang, Ye, and Xu 2017)	Inner (update)	$\mathcal{O}\left(\frac{m\cdot\log(n)}{n^2\epsilon^2}\right)$
DP-SCD	Inner (update)	$\mathcal{O}\left(\frac{L^3\cdot\log(\frac{n}{L})}{n^4\epsilon^2}\right)$

Notation

X Input dataset ($\mathbb{R}^{m \times n}$)
w Shared vector
α Dual vector
M Coordinate update mechanism
ε Privacy loss bound
N Gaussian noise (0, σ)
L mini-batch size
C scaling factor

Evaluation

Ridge regression

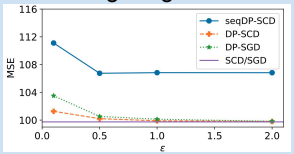

SVMs

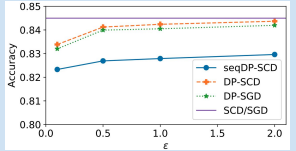

Logistic regression

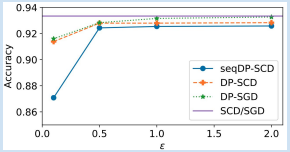

DP-SCD outperforms DP-SGD for applications that enable exact update steps (ridge regression and SVMs)

parallelism - scaling interplay

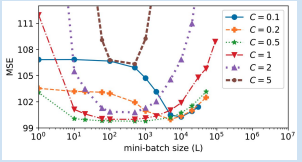

Deviating from the best choice for C (here: C_{best} = 0.5), reduces the width of the flat area and moves the minimum to the right (for smaller C values) or upwards (for larger C values)

https://github.com/gdamaskinos/dpscd

contact: georgios.damaskinos@gmail.com