Mutual-Information-Private Online Gradient Descent Algorithm

Ruochi Zhang, Parv Venkitasubramaniam. Electrical and Computer Engineering, Lehigh University, USA. • 04/18/2018

LEHIGH UNIVERSITY

Introduction

- **Online learning**, is to make a sequence of accurate predictions given knowledge of the correct answer to previous prediction tasks.
- Online learning applications: **targeted advertising** and **online ranking**.
- The purpose of this work is to propose **a user driven privatization mechanism** that allows the learner to infer the desired trends and patterns without compromising an individual users privacy.

Motivation

- Although individuals are willing to share their data, they are not expecting the disclosure of identities.
- The goal of learning is to uncover **"relationships"** or **"trends"** from historical data, which might be possible to be separated from the **information of individual identities**.
- The adversary may observe the **input data** of online learning system.

Related works

- A differentially private OCO method [P. Jain'12]
- Privacy-preserving deep learning [R. Shokri'15]

Full Information Online Convex Optimization (OCO)

Consider an online learning system that receives a stream of functions $(f_1, f_2, \cdots, f_T)$ and each $f_t : \mathscr{S} \to \mathbb{R}$ is a convex cost function representing data from one individual. The system is required to output a sequence of parameter estimates $(w_1, w_2, \cdots, w_T)$ with $w_t \in \mathscr{S} \subset \mathbb{R}^d$ that minimizes the total errors $\sum_{t=1}^{T} f_t(w_t)$. Due to causality, for every t, the algorithm computes w_t based only on $(f_1, f_2, \cdots, f_{t-1})$. We seek an algorithm $\mathscr{A}$ that minimize the **regret** defined by

$$\text{Regret}_T(\mathscr{A}) = \sum_{t=1}^{T} f_t(w_t) - \min_{w \in \mathscr{S}} \sum_{t=1}^{T} f_t(w)$$

We consider situations where

- the input functions $(f_1, f_2, \cdots, f_{t-1})$ are L-Lipschitz continuous
- the hypothesis space $\mathscr{S}$ is bounded w.r.t. l^2-norm. Under these restrictions, the OCO problem can be solved by the online gradient descent (OGD) algorithm [W. Davidon'76].

Algorithm 1 (A Privacy-preserving OGD)

Encryption layer:

Receive w_t from the learner
Pick a sub-gradient $z_t \in \partial f_t(w_t)$
Output $\tilde{z}_t = z_t + v_t$ to the learner, where $v_t \sim \mathcal{N}(0, \sigma^2 I)$ i.i.d.

Learner:

Receive $\tilde{z}_t$ from the encryption layer
Update $\theta_{t+1} = \theta_t - \tilde{z}_t$, (initialize $\theta_1 = 0$)
Predict $w_{t+1} = \arg\min_{w \in \mathscr{S}} \| w - \eta\theta_{t+1} \|$

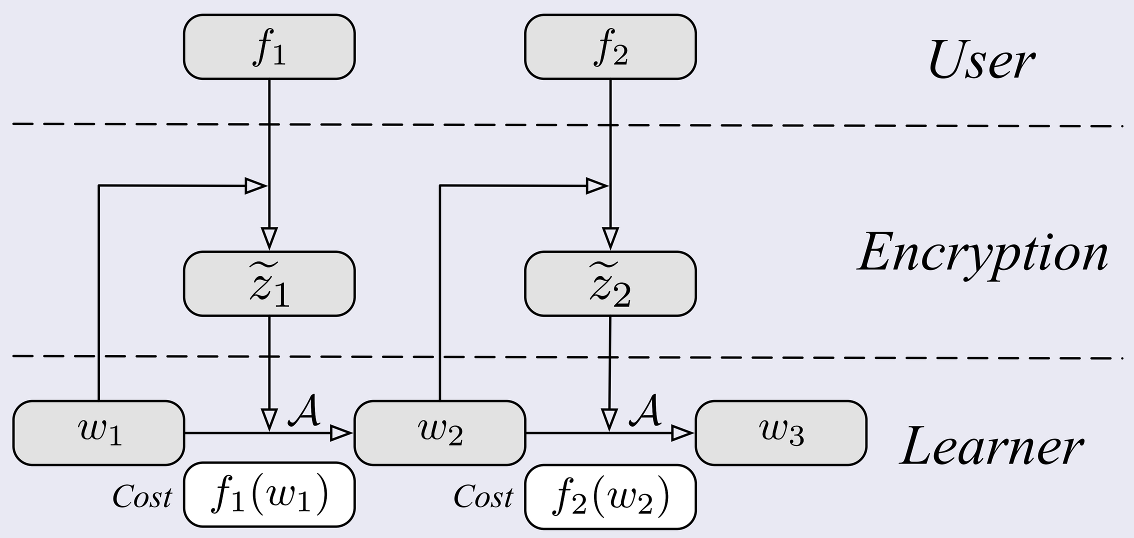

Theorem 1 (Privacy Guarantee)

The noise adding mechanism in Algorithm 1 is C-mutual-information private. i.e., $I(f_k; \tilde{z}_k) < C$ for every k, where $C = \frac{d}{2}\log(1 + \frac{L^2}{d\sigma^2})$

Theorem 2 (Regret Guarantee)

The Regret of Algorithm 1 is sub-linear to T. Specifically, $\text{Regret}(\mathscr{A}_1) \leq \frac{B^2}{2\eta} + \frac{\eta}{2}T(L^2 + d\sigma^2)$. In particular, by setting $\eta = B/\sqrt{(L^2 + d\sigma^2)T}$ we obtain the bound $\text{Regret}(\mathscr{A}_1) \leq B\sqrt{(L^2 + d\sigma^2)T}$.

Numerical Results (Full Information)

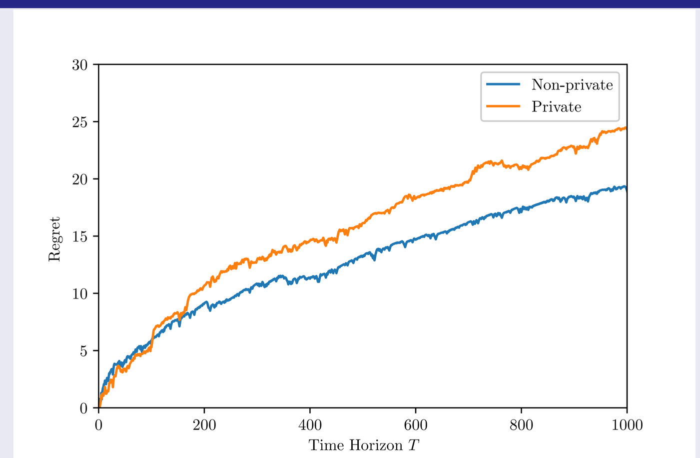

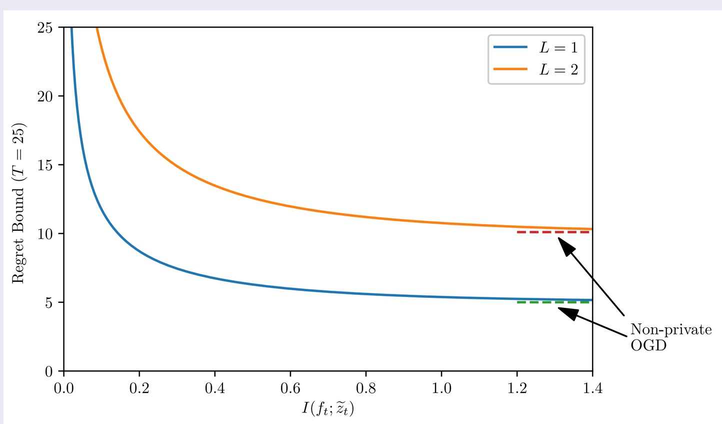

Extension: Bandit Setting OCO

Bandit setting [A. Flaxman'05]: for every t, the algorithm computes w_t based only on $(f_1(w_1), f_2(w_2), \cdots, f_{t-1}(w_{t-1}))$. i.e. The learner only knows the value of the loss function but he doesn't know the value of the loss function at other points.

Acknowledgement

National Science Foundation, grants CCF-1149495 and CCF-1617889.

Algorithm 2 (A Privacy-preserving OGD - Bandit)

Encryption layer:

Receive w_t from the learner
Pick $e_t \sim U_{sp}$, where U_{sp} is the uniform distribution over the unit sphere $\{u : \|u\|_2^2 = 1\}$.
Send $w_t + \delta e_t$ to the user
Receive cost value $\phi_t = f_t(w_t + \delta e_t)$ from the user
$z_t = \frac{d}{\delta}\phi_t e_t$
Output $\tilde{z}_t = z_t + v_t$ to the learner, where $v_t \sim \mathcal{N}(0, \sigma^2 I)$ i.i.d.

Learner:

Receive $\tilde{z}_t$ from the encryption layer
Update $\theta_{t+1} = \theta_t - \tilde{z}_t$, (initialize $\theta_1 = 0$)
Predict $w_{t+1} = \arg\min_{w \in \mathscr{S}} \| w - \eta\theta_{t+1} \|$

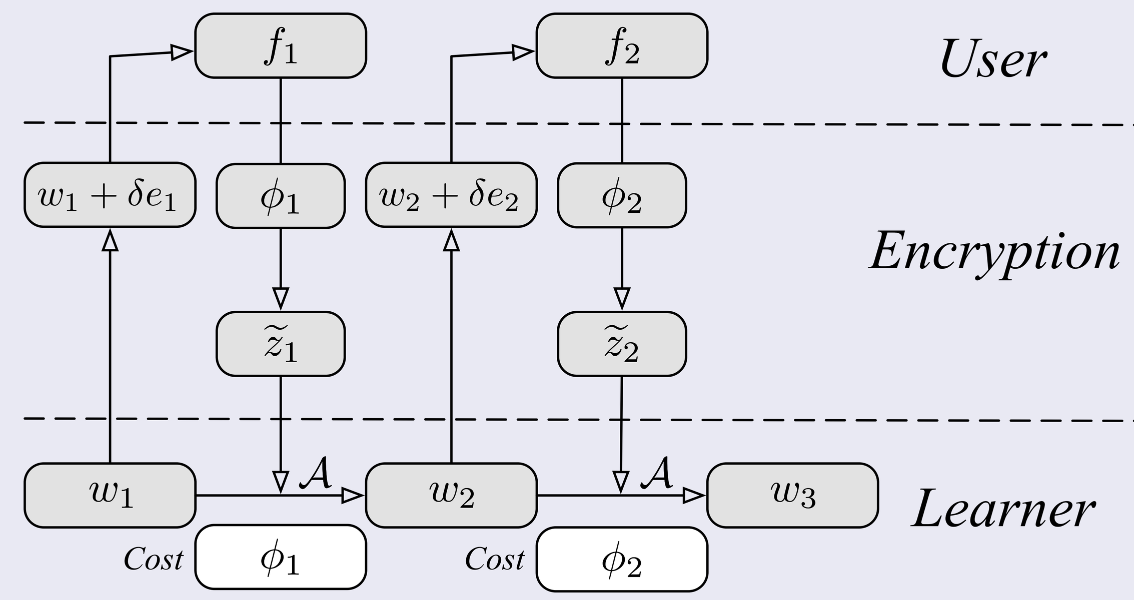

Theorem 3 (Privacy Guarantee - Bandit Setting)

Let $F = \max_{u \in \mathscr{S}, t \geq 1} f_t(u)$. If $F < \infty$, the proposed private OGD algorithm is C-mutual information private. i.e., $I(f_t; \tilde{z}_t) < C$ for every t, where $C = \frac{d}{2}\log(1 + \frac{d(F/\delta + L)^2}{\sigma^2})$

Theorem 4 (Regret Guarantee - Bandit Setting)

The Regret of the proposed private OGD algorithm is sub-linear to T. Specifically,

$$\text{Regret}(\mathscr{A}_2) \leq \frac{B^2}{2\eta} + \frac{\eta}{2}T(d^2(F/\eta + L)^2 + d\sigma^2) + 3TL\delta$$

In particular, if we set $\eta \sim T^{-3/4}$ and $\delta \sim T^{-1/4}$, the regret is bounded by $O(T^{3/4})$.

Conclusion

- Our private preserving OGD provides a conservative way to protect users' data.
- The user's leaked information is bounded by the channel capacity of Gaussian channel while the regret of the learning system is sub-linear to the time horizon T.