User-Specified Local Differential Privacy in Unconstrained Adaptive Online Learning

Dirk van der Hoeven

Universiteit Leiden
The Netherlands

Abstract

Standard Local Differential Privacy:
- All uses have the same privacy requirements.
- Learner chooses how much privacy users have.

User-Specified Local Differential Privacy:
- All uses have the same privacy requirements.
- Learner chooses how much privacy users have.

Adaptive Unconstrained algorithms:
- Compete with $\boldsymbol{u} \in \mathbb{R}^d$.
- Adapt to unknown and varying noise.
- Data dependent bounds.

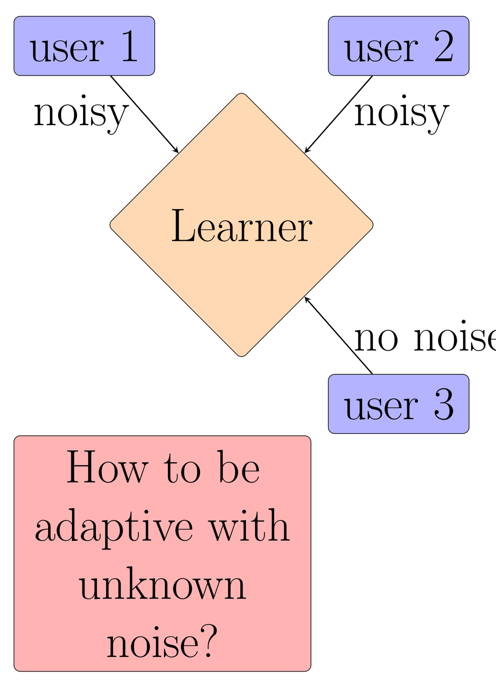

Private Unconstrained Online Convex Optimization with Noisy Subgradients

Local Differential Privacy: *Let $A = (X_1, \ldots, X_T)$ be a sensitive dataset where each $X_t \in A$ corresponds to data about individual t. A randomiser R which outputs a* ***disguised version*** *of $S = (U_1, \ldots, U_T)$ of A is said to provide ϵ-local differential privacy to individual t, if for all $x, x' \in A$ and for all $S \subseteq \mathcal{S}$,*

$$\Pr(U_t \in S | X_t = x) \leq \exp(\epsilon) \Pr(U_t \in S | X_t = x').$$

User specified LDP Unconstrained OCO:

Unconstrained comparator

0. **for** $t = 1, 2, \ldots, T$ **do**
1. The learner sends $\boldsymbol{w}_t \in \mathbb{R}^d$ to the provider of the t^{th} subgradient.
2. **The provider chooses** zero-mean and **symmetrical** ρ_t and samples $\boldsymbol{\xi}_t \sim \rho_t$.
3. The provider computes subgradient $\boldsymbol{g}_t \in \partial \ell_t(\boldsymbol{w}_t)$, where $\|\tilde{\boldsymbol{g}}_t\|_\star \leq G$.
4. The provider sends $\tilde{\boldsymbol{g}}_t = \boldsymbol{g}_t + \boldsymbol{\xi}_t \in \mathbb{R}^d$ to the learner.

Only bounded in $\mathbb{E}$ **Unknown to the learner**

Objective: minimize expected **regret** w.r.t. oracle parameter $\boldsymbol{u} \in \mathbb{R}$

$$E[\mathcal{R}_T(\boldsymbol{u})] = \sum_{t=1}^T \mathbb{E}[\langle \boldsymbol{w}_t, \tilde{\boldsymbol{g}}_t \rangle - \sum_{t=1}^T \langle \boldsymbol{u}, \tilde{\boldsymbol{g}}_t \rangle].$$

Goal: achieve adaptive expected regret bounds with unknown LDP requirements

$$\mathbb{E}[\mathcal{R}_T(\boldsymbol{u})] = O\left(\mathbb{E}[\|\boldsymbol{u}\| \sqrt{\sum_{t=1}^T \|\tilde{\boldsymbol{g}}_t\|_\star^2 \ln(1 + \|\boldsymbol{u}\|T))}]\right).$$

Previous work with known ρ_t:

$$O(\|\boldsymbol{u}\| \sqrt{(G^2 + \max_t \mathbb{E}[\|\boldsymbol{\xi}_t\|^2])T \ln(1 + \|\boldsymbol{u}\|T)})$$

Approach to Unconstrained OCO: Reward Regret Duality

Lemma 1. *If $-\mathbb{E}[\sum_{t=1}^T \langle \boldsymbol{w}_t, \boldsymbol{g}_t \rangle] \geq \mathbb{E}[F_T(-\sum_{t=1}^T \tilde{\boldsymbol{g}}_t) - c_T]$ for some convex function F_T and $c_T \in \mathbb{R}$, then $\mathbb{E}[\mathcal{R}_T(\boldsymbol{u})] \leq E[c_T] + F_T^\star(\boldsymbol{u})$.*

The difficulty lies in **finding a suitable F_T and c_T**. A route for noisy subgradients was provided by Jun and Orabona 2019: find an F_t, F_{t-1}, and $\boldsymbol{w}_t$ that satisfy

$$F_{t-1}(\boldsymbol{x}) - \langle \boldsymbol{w}_t, \boldsymbol{g}_t \rangle \geq \mathop{\mathbb{E}}_{\tilde{\boldsymbol{g}}_t}[F_t(\boldsymbol{x} - \tilde{\boldsymbol{g}}_t)], \tag{1}$$

then

$$\mathbb{E}[F_T(-\sum_{t=1}^T \tilde{\boldsymbol{g}}_t)] \leq \mathbb{E}[F_{T-1}(-\sum_{t=1}^{T-1} \tilde{\boldsymbol{g}}_t) - \langle \boldsymbol{w}_T, \boldsymbol{g}_T \rangle] \leq \ldots \leq \mathbb{E}[F_0(0) - \sum_{t=1}^T \langle \boldsymbol{w}_t, \boldsymbol{g}_t \rangle]$$

gives us what we need to use Lemma 1!

1d Potential Function

Key inequality: suppose $\boldsymbol{x}$ is a **symmetrical** random variable with $|\mathbb{E}[\langle \boldsymbol{v}, \boldsymbol{x} \rangle]| \leq \frac{1}{5}$ for some $\boldsymbol{v}$. Then

$$\mathbb{E}[\exp(\langle \boldsymbol{v}, \boldsymbol{x} \rangle - \langle \boldsymbol{v}, \boldsymbol{x} \rangle^2)] \leq 1 + \mathbb{E}[\langle \boldsymbol{v}, \boldsymbol{x} \rangle]. \tag{2}$$

Using (2) we can show that the following potential and predictions satisfy (1):
We use the 1d Potential function with (improper) prior P on $v \in [-\frac{1}{5G}, \frac{1}{5G}]$:

$$\mathbb{E}[F_t(-\sum_{s=1}^t \tilde{g}_s)] = \mathbb{E}[\mathop{\mathbb{E}}_{v \sim P}[\exp(-\sum_{s=1}^t v\tilde{g}_s - (v\tilde{g}_s)^2) - 1]].$$

Predictions:

$$w_t = \mathop{\mathbb{E}}_{v \sim P}[v \exp(-\sum_{s=1}^{t-1} v\tilde{g}_s - (v\tilde{g}_s)^2)]. \tag{3}$$

The proper priors we consider are of the form:

$$\frac{dP(v)}{dv} \propto \nu(v) \exp(-bv^2),$$

and improper prior:

$$\frac{dP(v)}{dv} = \frac{1}{|v|}$$

The predictions with prior $\frac{dP(v)}{dv} \propto \exp(-bv^2)$ satisfy:

$$\mathbb{E}[\mathcal{R}_T(u)] \leq 1 + |u| \max\left\{11G\left(\ln(|u|11G) - 1 + \ln\left(\frac{\sqrt{5}G\sqrt{\pi}}{4\sqrt{b}}\right)\right),\right.$$
$$\left.\mathbb{E}\left[\sqrt{8\left(b + \sum_{t=1}^T \tilde{g}_t^2\right) \ln(16|u|^2 \left(b + \sum_{t=1}^T \tilde{g}_t^2\right)^{\frac{3}{2}} \frac{\sqrt{\pi}}{\sqrt{b}} + 1)}\right]\right\}.$$

A More Granular Randomiser

For **more control** over the privacy we employ the Local Laplace randomiser:

$$\rho_t(\boldsymbol{z}) \propto \exp(-\sum_{j=1}^d \frac{\tau_{t,j}}{2} |\boldsymbol{z}_j|), \text{ where, } \sum_{j=1}^d \tau_{t,j} = \epsilon_t, \quad \tau_{t,j} \geq 0.$$

Application: Sparse Gradients

By running a copy of (3) in each dimension we obtain coordinate-wise we obtain an algorithm with AdaGrad type guarantees:

$$\mathbb{E}[\mathcal{R}_T(u)] \leq d + \sum_{j=1}^d |\boldsymbol{u}_j| \max\left\{11G\left(\ln(|\boldsymbol{u}_j|11G) - 1 + \ln\left(\frac{\sqrt{5}G\sqrt{\pi}}{4\sqrt{b_j}}\right)\right),\right.$$
$$\left.\mathbb{E}\left[\sqrt{8\left(b_j + \sum_{t=1}^T \tilde{g}_{t,j}^2\right) \ln(16|\boldsymbol{u}_j|^2 \left(b_j + \sum_{t=1}^T \tilde{g}_{t,j}^2\right)^{\frac{3}{2}} \frac{\sqrt{\pi}}{\sqrt{b_j}} + 1))}\right]\right\}.$$

In conjunction with the local laplace randomiser this could lead to smaller regret bounds if the privacy constraints match the sparsity.

Application: 1d to d Black-Box Reduction

We play $v_t \boldsymbol{w}_t$, where v_t is the prediction in (3) and $\|\boldsymbol{w}_t\| \leq 1$ comes from a standard OCO algorithm. By using the black-box reduction of Cutkosky and Orabona 2018 we obtain:

$$\mathbb{E}[\mathcal{R}_T(\boldsymbol{u})] = \underbrace{\mathbb{E}\left[\sum_{t=1}^T (v_t - \|\boldsymbol{u}\|)\langle \boldsymbol{z}_t, \tilde{\boldsymbol{g}}_t \rangle\right]}_{\text{regret for learning} \|\boldsymbol{u}\|} + \underbrace{\mathbb{E}\left[\sum_{t=1}^T \langle \boldsymbol{z}_t - \frac{\boldsymbol{u}}{\|\boldsymbol{u}\|}, \tilde{\boldsymbol{g}}_t \rangle\right]}_{\text{regret for learning} \frac{\boldsymbol{u}}{\|\boldsymbol{u}\|}}$$
$$= O\left(\mathbb{E}[\|\boldsymbol{u}\| \sqrt{\sum_{t=1}^T \|\tilde{\boldsymbol{g}}_t\|_\star^2 \ln(1 + \|\boldsymbol{u}\|T))}]\right)$$

Application: Private Adaptive Experts

Private expert setting:
1. Learner plays probability distribution $\boldsymbol{p}_t$ over experts.
2. Experts sample $\boldsymbol{\xi}_t \sim \rho_t$
3. Experts receive their losses $\boldsymbol{\ell}_t$.
4. Experts send their perturbed losses $\tilde{\boldsymbol{\ell}}_t = \boldsymbol{\ell}_t + \boldsymbol{\xi}_t$ to the learner.

The analysis of recent adaptive **expert algorithms** is **closely related** to the analysis in this paper. Using this relationship we can easily derive private and adaptive expert algorithms. Replace $\tilde{g}_t$ by the instantaneous regret $\tilde{r}_t(i) = \langle \boldsymbol{p}_t - \boldsymbol{e}_i, \tilde{l}_t \rangle$, where $\tilde{\boldsymbol{\ell}}_t = \boldsymbol{\ell}_t + \boldsymbol{\xi}_t$ and $\boldsymbol{\ell}_t$ is a vector containing the expert losses. If we play

$$p_t(i) \propto p_1(i) \mathop{\mathbb{E}}_{v \sim P}[v \exp(-\sum_{s=1}^{t-1} v\tilde{r}_t(i) - (v\tilde{r}_t(i))^2)],$$

where P is a prior on $v \in [0, \frac{1}{5G}]$
It can be shown that, by using (2) instead of the prod bound in the analysis of squint, that the expected regret satisfies:

$$\mathbb{E}[\mathcal{R}_T(i)] = O\left(\mathbb{E}\left[\sqrt{\ln(k) \sum_{t=1}^T \tilde{r}_t(i)^2}\right]\right)$$

LaTeX TikZposter