

Private learning / Oprim

work in progress with.

Etienne Boursier

ENS Paris-Saclay

and Criteo AI Lab

Privacy?

* differential privacy:
"cannot retrieve specific info
in databases"

* private learning:
"use/exploit some
private info about you
without revealing too much"

Privacy for users

- * I want to visit websites I like... without Google knowing exactly how depraved I am.
- * I want to watch movies on Netflix without being classified as white/male/30's.

Film fans see red over Netflix 'targeted' posters for black viewers

The streaming service's customers say they are being duped by marketing that shows minor cast members as leading characters



▲ Set It Up is made to look like a two-hander between Taye Diggs and Lucy Liu, rather than the white couple.
Photograph: Twitter Kelly Quantrill @codetrill

Privacy for a company

- * Start-up wants to use Deep Learning / High Perf Comp.
 - Send data to Amazon
Microsoft
IBM ...
 - They do the computations.
 - Start-up gets the result.

Another Concrete Example

- * Criteo is buying "ad slots" via auctions
- * These auctions are run by Google
- * The major competitor of Criteo is Google

A simple model

- $\text{Max}_{x \in X \subset \mathbb{R}^d} x^T A_k y$ $y \in \mathbb{R}^n$
 $A_k \in \mathbb{R}^{d \times n}$
- $k \in \{1, \dots, K\}$ is the private type
prior $\pi_0 \in \Delta_K$, $k \sim \pi_0$.
- Privacy value is the amount of info leaked on k .
ex $KL(\pi_0, \pi_1)$

What is the posterior π_1 ?

$$\max_{x \in \mathcal{X}} x^T A_k y, \quad k \sim \pi_0$$

• Given $k \in \{1, \dots, K\}$, choose $x \sim \mu_k \in \mathcal{P}(\mathcal{X})$

• $\pi_1 \in \Delta_K$ posterior knowing x

$$\pi_1^k | x = \frac{\pi_0^k \mu_k(x)}{\sum_{j=1}^K \pi_0^j \mu_j(x)}, \quad \text{Bayes}$$

Private learning obj.

$$\max_{P_1, \dots, P_K} \sum_{k=1}^K \pi_0^k \mathbb{E}_{x \sim P_k} [x^T A_k y - \lambda k L(\pi_{|x}, \pi_0)]$$

more generally

$$\inf_{\gamma \in \mathcal{P}(X \times [K])} \int c(x, k) + \lambda \mathcal{D}(\pi_x, \pi_0) d\gamma(x, k)$$

$\pi_{\sharp} \gamma = \pi_0$

The f -divergence case

• f convex, $f(1) = 0$

$$\mathcal{D}(P, Q) = \mathbb{E}_{x \sim Q} \left[f \left(\frac{P(x)}{Q(x)} \right) \right]$$

ex $KL(Q, P) = \mathbb{E}_{x \sim Q} \left[-\log \left(\frac{P(x)}{Q(x)} \right) \right]$

$$KL(P, Q) = \mathbb{E}_{x \sim Q} \left[\frac{P(x)}{Q(x)} \log \left(\frac{P(x)}{Q(x)} \right) \right]$$

$$TV(P, Q) = \mathbb{E}_{x \sim Q} \left[\frac{1}{2} \left| \frac{P(x)}{Q(x)} - 1 \right| \right]$$

Convexity result

f convex, $f(1) = 0$

$$\inf_{\substack{\gamma \in \Delta(X_k) \\ \pi_2 \# \gamma = \pi_0}} \left(\int c d\gamma + \lambda \int \mathbb{E} f\left(\frac{d\pi_{12}}{d\pi_0}\right) d\gamma \right)$$

is convex in γ !

Finiteness result

If k is finite (card k).

Th. $\forall \varepsilon > 0, \exists \varepsilon$ -optimal γ with
finite support $k(k+2)$

• X compact + $c(\cdot, k)$ lsc
true for $\varepsilon = 0$

Finite reformulation?

$$\inf_f \sum_{i,k} \gamma_{i,k} c(x^i, k) + \lambda \sum_{i,k,j} \gamma_{i,k} \pi_{i,j}^j f\left(\frac{\gamma_{i,j}}{\pi_{i,j}^j \sum_p \gamma_{i,p}}\right)$$

$$\left[\begin{array}{l} \gamma \in \mathbb{R}^{(k+2)k} \\ x \in \mathbb{R}^{k+2} \\ \sum_i \gamma_{i,j} = \pi_{i,j}^j \end{array} \right]$$

No longer convex!

Ex: the linear case

$$c(x, k) = w_k^T x + \beta_k, \quad x \in [-1, 1]^d$$

$$\inf_{\gamma} \sum_i \left\| \sum_k \gamma_{i,k} w_k \right\|_1 + \sum_k \pi_0^k \beta_k + \lambda \sum_{i,k,j} \gamma_{i,k} \pi_0^j \log \left(\frac{\gamma_{i,j}}{\pi_0^j \sum_p \gamma_{i,p}} \right)$$

$$= -G(\gamma) + F(\gamma)$$

→ Difference of convex

KL-divergence and OT

$$\inf_{\substack{\gamma \in \mathcal{P}(X \times [k]) \\ \pi_{\#} \gamma = \nu}} \int c(x, k) + \lambda \mathbb{E} \left| \log \left(\frac{d\pi_x}{d\nu} \right) \right| d\gamma(x, k)$$

||

(ν is the prior)

$$\inf_{\mu \in \mathcal{P}(X)} \inf_{\pi \in \mathcal{T}(\mu, \nu)} \int c d\pi + \lambda \int \log \left(\frac{d\pi}{d\mu d\nu} \right) d\pi$$

OT_{c, λ}(μ, ν)

Regularized OT

$$\mathcal{O}T_{c,1}(\mu, \nu) = \inf_{\pi \in T(\mu, \nu)} \int c d\pi + \lambda \int \log\left(\frac{d\pi}{d\mu d\nu}\right) d\pi$$

$\hookrightarrow$ solve with Sinkhorn (in π)

$\hookrightarrow$ Optimize (in μ)?

Parametric family

$$\min_{\mu \in \mathcal{P}(x)} OT_{c, \lambda}(\mu, \nu)$$

↳ look for $\mu_{\theta} = \sum_{j=1}^n \alpha_j(\theta) \delta_{x_j(\theta)}$

Compute $\frac{\partial}{\partial \alpha} OT_{c, \lambda}(\mu, \nu)$ and $\frac{\partial}{\partial x} OT_{c, \lambda}(\mu, \nu)$

(automatic diff - dual)

Applications

Private LP: $c(x, k) = x^T w_k$,

$$v = \sum \pi_0^k \delta_{w_k}$$

Private Regression: $v = \sum \pi_0^k \delta_{A_k}$

$$c(A_k, A) = \left\| (X^T A^T A X)^{-1} - (X^T A_k^T A_k X)^{-1} \right\|_2$$

$$(y = A_k X \beta + \varepsilon; (y, x) \text{ public } A_k \text{ private})$$

So many questions...

- * Statistical guarantees
- * Computational issues
- * Private optimization algo
(query AWS repeatedly)
- * General f -divergence ...