Utility/Privacy Trade-off through the lens of Optimal Transport

Etienne Boursier [1]   Vianney Perchet [2, 3]

[1] ENS Paris-Saclay, CMLA

[2]Criteo AI Lab, Paris

[3]ENSAE Paris

MLITRW '19, Criteo Paris

An economic motivation

Online repeated auctions

Ad slot valued v. Bid $p \implies$ auctioneer infers v.
Auctioneer's revenue $\nearrow$ while bidder's utility $\searrow$ when v public.

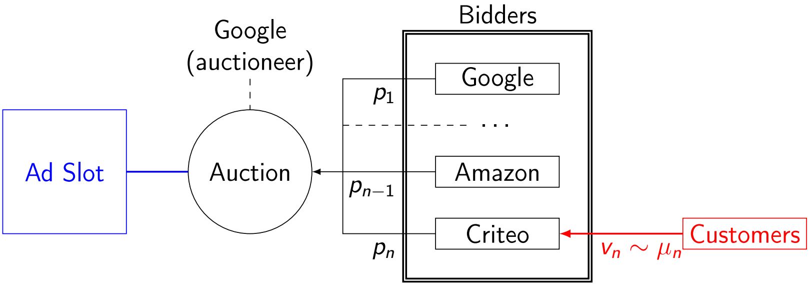

Online advertisement auction system

Bidder's goal: short term utility **and** hide value distribution μ_n

Toy example

Player: minimizes utility loss

$$\min_{x \in \mathcal{X} \subset \mathbb{R}^d} x^\top y_k$$

y_k depends on **private type** $k \in \{1, \ldots, K\}$ with prior $p_0 \in \Delta_K$.
Adversary: observes x and infers k

Program in previous literature[1]:

$$\min_{\mu_1, \ldots, \mu_K} \sum_{k=1}^{K} p_0(k) \, \mathbb{E}_{x \sim \mu_k} \left[x^\top y_k \right]$$
$$\text{such that } \mathbb{E}[KL(p_x, p_0)] \leq \varepsilon$$

[1] Eilat, R., Eliaz, K., and Mu, X. (2019). Optimal privacy-constrained mechanisms

General formulation of the problem

Our general program:

$$\inf_{\substack{\gamma \in \mathcal{P}(\mathcal{X}\times\mathcal{Y}) \\ \pi_2\#\gamma = p_0}} \int_{\mathcal{X}\times\mathcal{Y}} \left(c(x,y) + \lambda\, D(p_x, p_0)\right) \mathrm{d}\gamma(x,y) \quad \text{(P-OPT)}$$

- type $y \sim p_0 \in \mathcal{P}(\mathcal{Y})$
- $\pi_2\#\gamma(A) = \gamma(\mathcal{X} \times A)$
- $c =$ utility loss ; $D =$ privacy loss (e.g. KL)

Theoretical results

Theorem (Convexity)

If D is an f-divergence, then (P-OPT) *is convex in γ.*

$\rightarrow$ (P-OPT) easy for finite $\mathcal{X}$ and $\mathcal{Y}$.

Theorem (Finite prior support)

If $|\mathrm{supp}(p_0)| = K$, for all $\varepsilon > 0$, we can look for a solution of (P-OPT) *with support of size $K(K+2)$.*

$\rightarrow$ finite dimension ☺ but not jointly convex ☹

Sinkhorn divergence minimization

Definition (Sinkhorn divergence)

$$\mathrm{OT}_{c,\lambda}(\mu,\nu) = \inf_{\gamma \in \Pi(\mu,\nu)} \int c \mathrm{d}\gamma + \lambda \int \log\left(\frac{\mathrm{d}\gamma}{\mathrm{d}\mu \mathrm{d}\nu}\right) \mathrm{d}\gamma$$

- entropic regularization $\implies$ fast OT distances approximation[2]

If D=KL, (P-OPT) equivalent to

$$\inf_{\mu \in \mathcal{P}(\mathcal{X})} \mathrm{OT}_{c,\lambda}(\mu, p_0).$$

[2] Cuturi, M. (2013). Sinkhorn distances: Lightspeed computation of optimal transport

Recap

- utility-privacy trade-off motivated by economic mechanisms
- general regularized problem
- convexity + finiteness under mild assumptions
- benefit from Sinkhorn divergence
- find our simulations in the paper

Slides, code and paper at `eboursier.github.io`

Thank you !