Big idea 4: Optimality and Fundamental Limits

John Duchi
Stanford University

Frejus 2025

Outline for today

1. optimality techniques in local differential privacy
2. optimality techniques in central differential privacy

Setting for lower bounds

Definition (Minimax risk)

For parameter $\theta = \theta(P)$ of interest, *minimax risk* for the loss ℓ is

$$\inf_{M \in \mathcal{M}} \sup_{P \in \mathcal{P}} \mathbb{E}\left[\ell(M(P_n) - \theta(P))\right]$$

where infimum is over family $\mathcal{M}$ of mechanisms

Basic lower bound techniques: from estimation to testing

- call a pair P_0, P_1 of distributions δ-separated if

$$|\theta(P_0) - \theta(P_1)| \geq \delta$$

Lemma (Le Cam's method; cf. [6])

For any two distributions P_0 *and* P_1,

$$\begin{aligned} \max_{P \in \{P_0, P_1\}} \mathbb{E}_P \left[\ell \left(|\widehat{\theta} - \theta(P)| \right) \right] &\geq \frac{\ell(\delta/2)}{2} \inf_{\Psi} \left\{ P_0(\Psi = 1) + P_1(\Psi = 0) \right\} \\ &= \frac{\ell(\delta/2)}{2} \left(1 - \|P_0 - P_1\|_{\mathrm{TV}} \right). \end{aligned}$$

Example lower bound technique

Big idea: find parameters as far apart as possible while hard to test between P_0, P_1 (see Duchi [6] for more)

Proposition (Pinsker's inequality)

For distributions P_0, P_1,

$$\|P_0 - P_1\|_{\mathrm{TV}}^2 \le \frac{1}{2} D_{\mathrm{kl}}\left(P_0 \| P_1\right)$$

smaller idea: often $D_{\mathrm{kl}}\left(P_0 \| P_1\right) \lesssim \delta^2$ for $|\theta(P_0) - \theta(P_1)| \le \delta$

Example (distance between normals)

For $P_0 = \mathsf{N}(\mu_0, \sigma^2)$ and $P_1 = \mathsf{N}(\mu_1, \sigma^2)$,

$$D_{\mathrm{kl}}\left(P_0 \| P_1\right) = \frac{(\mu_0 - \mu_1)^2}{2\sigma^2}$$

Example lower bound technique (continued)

for δ-separated P_0, P_1,

$$\max_{P \in \{P_0, P_1\}} \mathbb{E}_{P^n}\left[\ell\left(|\widehat{\theta} - \theta(P)|\right)\right] \geq \frac{\ell(\delta/2)}{2}\left(1 - \|P_0^n - P_1^n\|_{\mathrm{TV}}\right)$$

- in "typical" case that $D_{\mathrm{kl}}\left(P_0 \| P_1\right) \leq \kappa\delta^2$ for $|\theta(P_0) - \theta(P_1)| = \delta$,

$$\|P_0^n - P_1^n\|_{\mathrm{TV}}^2 \leq \frac{1}{2} D_{\mathrm{kl}}\left(P_0^n \| P_1^n\right) = \frac{n}{2} D_{\mathrm{kl}}\left(P_0 \| P_1\right) \leq \frac{\kappa n \delta^2}{2}$$

- make probability of error $\frac{1}{2}$:

$$\delta^2 = \frac{1}{2\kappa n}$$

- lower bound

$$\frac{1}{4}\ell\left(\frac{1}{2\sqrt{2\kappa n}}\right).$$

Example (Normal estimation lower bound)

For location estimation in $\{\mathsf{N}(\theta, \sigma^2)\}_{\theta \in \mathbb{R}}$,

$$\sup_P \mathbb{E}_{P^n}\left[\ell\left(|\widehat{\theta}_n - \theta(P)|\right)\right] \geq \frac{1}{4}\ell\left(\frac{\sigma}{2\sqrt{n}}\right)$$

Lower bound in locally private scenarios

- data release via (sequentially interactive) channel:

$$Z_i \sim Q(\cdot \mid X_i, Z_1^{i-1})$$

  where Q is ε-differentially private

- interested in *locally private* minimax risk

$$\mathfrak{M}_n(\varepsilon) := \inf_{Q_1^n} \inf_{\widehat{\theta}} \sup_{P \in \mathcal{P}} \mathbb{E}\left[\ell\left(\widehat{\theta}(Z_1^n) - \theta(P)\right)\right]$$

The key contraction

For distributions P_0, P_1, let R_v^n be the *result* marginals over Z_1^n from

$$X_i \stackrel{\text{iid}}{\sim} P_v, \quad Z_i \sim Q(\cdot X_i, Z_1^{i-1})$$

Theorem (Duchi et al. [9], Corollary 3)

For any sequentially interactive ε-locally differentially private channels,

$$D_{\text{kl}}\left(R_0^n \| R_1^n\right) \leq 4n(e^{\varepsilon} - 1)^2 \left\| P_0 - P_1 \right\|_{\text{TV}}^2 .$$

Generic lower bounds

Corollary

For any pair of distributions with $|\theta(P_0) - \theta(P_1)| \geq \delta$*,*

$$\mathfrak{M}_n(\varepsilon) \gtrsim \ell(\delta/2)\left(1 - 4n(e^{\varepsilon} - 1)^2 \|P_0 - P_1\|_{\mathrm{TV}}^2\right).$$

Example (Mean estimation with k moments)

If $\mathcal{P}_k = \{P : \mathbb{E}_P[|X|^k] \leq 1\}$, then minimax mean-squared error has scaling

$$\mathfrak{M}_n(\varepsilon) \asymp \left(\frac{1}{n(e^{\varepsilon} - 1)^2}\right)^{\frac{k-1}{k}}.$$

Lower bounds in central differential privacy

Big picture: let d-dimensional estimator converges with rate $r(n)$, i.e.,

$$\mathbb{E}[\ell(\widehat{\theta}_n - \theta)] \asymp r(n)$$

with ε-differential privacy, expect privacy penalty

$$\mathbb{E}[\ell(\widehat{\theta}_n - \theta)] \asymp r(n) + r\left(\frac{n^2\varepsilon^2}{d^2}\right)$$

with $(\varepsilon, \delta$-differential privacy, expect penalty

$$\mathbb{E}[\ell(\widehat{\theta}_n - \theta)] \asymp r(n) + r\left(\frac{n^2\varepsilon^2}{d\log(1/\delta)}\right)$$

Example

mean estimation with data $x_i \in \mathbb{R}^d$, $\|x_i\|_2 \leq 1$

Cai et al.'s Score Attack

Definition

The *(Fisher) score* is

$$s_\theta(x) := \nabla \log p_\theta(x) = \frac{\nabla p_\theta}{p_\theta}(x)$$

Idea: (Cai et al. [4, 5]): if $M(P_n)$ is an accurate estimator, then

1. $M(P_n)$ should correlate with $\sum_{i=1}^n s_\theta(X_i)$, but
2. privacy limits this correlation

The minimaxlower bound

Theorem (Cai et al. [5])

Define Fisher Information $I_\theta := \mathbb{E}[s_\theta(X)s_\theta(X)^T]$ *and let* M *be* (ε, δ)*-differentially private. For any smooth enough prior* π *on* θ *near* θ_0*,*

$$\int \mathbb{E}_\theta \left[\|M(P_n) - \theta\|_2^2 \right] \pi(\theta) d\theta \gtrsim \frac{d^2}{n^2\varepsilon^2} \cdot \frac{1}{\|I_{\theta_0}\|_{\mathrm{op}}}.$$

Remark: classical lower bounds scale as $\frac{d}{n\|I_\theta\|_{\mathrm{op}}}$

Example (Gaussian mean estimation)

Let $X_i \stackrel{\mathrm{iid}}{\sim} \mathsf{N}(\theta, I_d)$, where $\|\theta\|_2 \leq 1$. Then

$$\int \mathbb{E}_\theta \left[\|M(P_n) - \theta\|_2^2 \right] \pi(\theta) d\theta \gtrsim \frac{d}{n} + \frac{d^2}{n^2\varepsilon^2}.$$

Proof I

Define *alignment*

$$A_\theta(x, P_n) := \langle M(P_n) - \theta, s_\theta(x) \rangle$$

and let $X' \sim P_\theta$, independent of P_n

Lemma

we have $\mathbb{E}[A_\theta(X', P_n)] = 0$ *and*

$$\mathbb{E}[|A_\theta(X', P_n)|] \le \sqrt{\mathbb{E}[\|M(P_n) - \theta\|_2^2]} \cdot \|I_\theta\|_{\mathrm{op}}^{1/2}$$

Proof II: bounding alignment by privacy

Lemma

We have $\mathbb{E}[A_\theta(X, P_n)] \leq (e^\varepsilon - 1)\mathbb{E}[|A_\theta(X', P_n)|]$.

Proof III: from alignment to expectations

Lemma

The summed alignment satisfies

$$\sum_{i=1}^{n} \mathbb{E}[A_\theta(X_i, P_n)] = \sum_{j=1}^{d} \frac{\partial}{\partial \theta_j} \mathbb{E}_\theta[M_j(P_n)]$$

Lemma (Proposition 2.2 [5])

If $\mathbb{E}[\|M(P_n) - \theta\|_2^2] = O(1)$, *then*

$$\sum_{j=1}^{d} \int \frac{\partial}{\partial \theta_j} \mathbb{E}_\theta[M_j(P_n)] \pi(\theta) d\theta \gtrsim d.$$

Putting it all together

$$
\begin{aligned}
d &\lesssim \sum_{i=1}^{n} \mathbb{E}[A_\theta(X_i, P_n)] \\
&\leq n(e^{\varepsilon} - 1)\mathbb{E}\left[|A_\theta(X', P_n)|\right] \\
&\leq n(e^{\varepsilon} - 1)\mathbb{E}\left[\|M(P_n) - \theta\|_2^2\right]^{1/2} \|I_\theta\|_{\mathrm{op}}^{1/2} .
\end{aligned}
$$

A few additional references

- Optimality in local differential privacy:
  - Duchi and Rogers [7] present general (interactive) lower bounds using communication complexity
  - Duchi and Ruan [8] present a "geometric" characterization of local differential privacy (asymptotics)
  - Acharya et al. [1] present results on information-constrained estimation
- Optimality in central differential privacy:
  - early work using pure differential privacy and packings [11, 14, 3]
  - Steinke and Ullman [15] leverage fingerprinting (cryptographic) lower bounds [10, 12, 13]
  - Attias et al. [2] provide lower bounds on memorization in statistical learning using similar "score attack" techniques

Take-homes

1. Many open questions remain in privacy
   - continuous observation (e.g., long-term users)
   - leveraging public data
   - fundamental limits
   - even basic statistical questions
2. Big ideas we've discussed
   - Definitions and importance of composition
   - Amplification: shuffling, sampling, iteration
   - Some more sophisticated mechanisms (inverse sensitivity, matrix mechanisms)
   - Optimality
3. Big ideas we've missed
   - propose-test-release framework
   - application areas and deployments, e.g., machine learning, US Census
   - others!

[1] J. Acharya, C. Cannone, C. Freitag, Z. Sun, and H. Tyagi. Inference under information constraints III: Local privacy constraints. *IEEE Journal on Selected Areas in Information Theory*, 2(1):253–267, 2021.

[2] I. Attias, G. K. Dziugaite, M. Haghifam, R. Livni, and D. M. Roy. Information complexity of stochastic convex optimization: Applications to generalization and memorization. In *Proceedings of the 41st International Conference on Machine Learning*, 2024.

[3] R. F. Barber and J. C. Duchi. Privacy and statistical risk: Formalisms and minimax bounds. *arXiv:1412.4451 [math.ST]*, 2014.

[4] T. T. Cai, Y. Wang, and L. Zhang. The cost of privacy: optimal rates of convergence for parameter estimation with differential privacy. *Annals of Statistics*, 49(5):2825–2850, 2021.

[5] T. T. Cai, Y. Wang, and L. Zhang. Score attack: A lower bound technique for optimal differentially private learning. *arXiv:2303.07152 [math.ST]*, 2023.

[6] J. C. Duchi. *Lecture Notes on Information Theory and Statistics*. 2024. URL `http://web.stanford.edu/class/stats311`.

[7] J. C. Duchi and R. Rogers. Lower bounds for locally private estimation via communication complexity. In *Proceedings of the*

Thirty Second Annual Conference on Computational Learning Theory, 2019.

[8] J. C. Duchi and F. Ruan. The right complexity measure in locally private estimation: It is not the Fisher information. *Annals of Statistics*, 52(1):1–51, 2024.

[9] J. C. Duchi, M. I. Jordan, and M. J. Wainwright. Minimax optimal procedures for locally private estimation (with discussion). *Journal of the American Statistical Association*, 113(521):182–215, 2018.

[10] C. Dwork, A. Smith, T. Steinke, J. Ullman, and S. Vadhan. Robust traceability from trace amounts. In *56th Annual Symposium on Foundations of Computer Science*, 2015. Long version available at `https://jonathan-ullman.github.io/assets/tracing.pdf`.

[11] M. Hardt and K. Talwar. On the geometry of differential privacy. In *Proceedings of the Forty-Second Annual ACM Symposium on the Theory of Computing*, pages 705–714, 2010.

[12] G. Kamath, J. Li, V. Singhal, and J. R. Ullman. Privately learning high-dimensional distributions. In *Proceedings of the Thirty Second Annual Conference on Computational Learning Theory*, 2019.

[13] G. Kamath, A. Mouzakis, and V. Singhal. New lower bounds for private estimation and a generalized fingerprinting lemma. In *Advances in Neural Information Processing Systems 35*, 2022.

[14] A. Nikolov, K. Talwar, and L. Zhang. The geometry of differential privacy: the sparse and approximate case. In *Proceedings of the Forty-Fifth Annual ACM Symposium on the Theory of Computing*, 2013.

[15] T. Steinke and J. Ullman. Between pure and approximate differential privacy. In *Proceedings of the Twenty Eighth Annual Conference on Computational Learning Theory*, 2015.