Big idea 2: privacy amplification

John Duchi
Stanford University

Frejus 2025

The plan

- subsampling
- stochastic gradient methods
- shuffling and local privacy

Intuition

- algorithmic choices ought to improve privacy
- subsampling: only a small fraction of individuals?

- local privacy and shuffling

The subsampling mechanism

- assume mechanisms M work on $\mathcal{P}_1, \mathcal{P}_2, \ldots, \mathcal{P}_n$
- the *subsampling mechanism*

$$\text{draw } S \subset \{1, \ldots, n\} \text{ with } \operatorname{card}(S) = m < n$$

uniformly at random, release

$$M^{\mathsf{sub}}(P_n) := M(P_S), \quad P_S = \frac{1}{m} \sum_{i \in S} \mathbf{1}_{X_i}$$

Privacy amplification of subsampling

Theorem (Ullman [12], Balle et al. [4])

Let $q = \frac{m}{n}$. If M is (ε, δ)-differentially private, then M^{sub} is

$$(\log(1 + q(e^{\varepsilon} - 1)), q\delta)\text{-DP.}$$

Corollary

If $\varepsilon \leq 1$, then M^{sub} is $(O(1)q\varepsilon, q\delta)$-differentially private

Proof of amplification I

- Let P_n, P'_n differ on index i, $q = \frac{m}{n}$, write

$$\begin{aligned}
&\mathbb{P}(M^{\mathsf{sub}}(P_n) \in A) \\
&= (1-q)\mathbb{P}(M(P_S) \in A \mid i \notin S) + q\mathbb{P}(M(P_S) \in A \mid i \in S)
\end{aligned}$$

Proof of amplification II

- bound the terms

$$p_0 = \mathbb{P}(M(P_S) \in A \mid i \notin S),$$
$$p_1 = \mathbb{P}(M(P_S) \in A \mid i \in S), \quad p_1' = \mathbb{P}(M(P_S') \in A \mid i \in S)$$

A heuristic application: mean estimation

- assume data $x_1, \dots, x_n \in [-1, 1]$,
- choose $\sigma^2_{\varepsilon,\delta} = O(1)\frac{\log \frac{1}{\delta}}{\varepsilon^2}$ so $x_i + \mathsf{N}(0, \sigma^2_{\varepsilon,\delta})$ is (ε, δ)-private
- repeat k times: sample $S_j \subset [n]$, $m = \mathrm{card}(S_j) = qn$,

$$\hat{\mu}_j = \frac{1}{m} \sum_{i \in S_j} x_i + \mathsf{N}\left(0, \frac{\sigma^2_{\varepsilon,\delta}}{m^2}\right).$$

Proposition (Informal)

If $q \geq \frac{1}{n}$ and $k \leq \frac{1}{q^2}$, then $\widehat{\mu} = \frac{1}{k} \sum_{j=1}^{k} \widehat{\mu}_j$ is (ε, δ)-private and

$$\mathbb{E}\left[\left(\widehat{\mu} - \overline{x}_n\right)^2\right] \leq \frac{1}{kqn} + \frac{\sigma^2_{\varepsilon,\delta}}{km^2}.$$

Idea of argument

Application of amplification: stochastic gradient

Problem: minimize

$$f(\theta) := \frac{1}{n} \sum_{i=1}^{n} f_i(\theta)$$

where f_i are convex, L_0-Lipschitz, have L_1-Lipschitz gradient

Example (Logistic regression)

For data $(x_i, y_i) \in \mathbb{R}^d \times \{-1, 1\}$, $\|x_i\|_2 \leq L_0$, losses

$$f_i(\theta) = \log(1 + \exp(-y_i x_i^T \theta))$$

are L_0-Lipschitz and have $L_1 = \frac{L_0}{4}$.

Application of amplification: stochastic gradient

Problem: minimize

$$f(\theta) := \frac{1}{n} \sum_{i=1}^{n} f_i(\theta)$$

Algorithm: noisy SGD

- sample $S \subset [n]$, $\mathrm{card}(S) = m$, compute

$$g_k = \frac{1}{m} \sum_{i \in S} \nabla f_i(\theta_k) + \mathsf{N}(0, \sigma^2 I)$$

- update $\theta_{k+1} = \theta_k - \eta_k g_k$

Convergence of private SGD

Theorem (Folklore, see [2], Thm. 1)

If $\mathrm{Var}(g_k \mid \theta_k) \leq V^2$, *then stepsize* $\eta = \frac{1}{L_1+\gamma}$

$$\sum_{i=1}^{k} \mathbb{E}[f(\theta_{i+1}) - f(\theta^\star)] \leq \frac{\|\theta_1 - \theta^\star\|^2}{2\eta} + \frac{k}{2\gamma} V^2.$$

Corollary (Noisy SGD)

$$\mathbb{E}[f(\overline{\theta}_k) - f(\theta^\star)] \leq \frac{(L_1+\gamma)\,\|\theta_1 - \theta^\star\|^2}{2k} + \frac{1}{2\gamma}\left(\sigma^2 d + \frac{L_0^2}{m}\right).$$

Privacy analysis of noisy SGD

- sensitivity of $g = \frac{1}{m} \sum_{i \in S} \nabla f_i(\theta_k)$:

$$\left\| g - g' \right\|_2 \leq \frac{2L_0}{m}$$

- to obtain (ε, δ)-DP per iteration via $g_k = g + \mathsf{N}(0, \sigma^2 I_d)$:

$$\sigma^2 = O(1) \frac{L_0^2}{m^2} \cdot \frac{\log \frac{1}{\delta}}{\varepsilon^2}.$$

- with subsampling $q = \frac{m}{n}$ and $\varepsilon \leq 1$, per-step privacy

$$\varepsilon_{\mathsf{step}} \lesssim q\varepsilon$$

Privacy analysis of noisy SGD (continued)

- total privacy loss after k iterations:

$$\left(k\varepsilon_{\text{step}}^2 + O(1)\sqrt{k\varepsilon_{\text{step}}^2 \log\frac{1}{\delta}}, k\varepsilon \right) \text{-differentially private.}$$

  i.e.

$$\varepsilon_{\text{total}} \lesssim kq^2\varepsilon^2 + \sqrt{k(q\varepsilon)^2 \log\frac{1}{\delta}}$$

- convergence (with $R^2 = \|\theta_1 - \theta^\star\|_2^2$)

$$\mathbb{E}[f(\overline{\theta}_k) - f(\theta^\star)] \lesssim \frac{L_1 R^2}{k} + \frac{\gamma R^2}{k} + \frac{1}{\gamma}\left(\frac{dL_0^2 \log\frac{1}{\delta}}{m^2\varepsilon^2} + \frac{L_0^2}{m} \right)$$

- find dominant terms

Convergence of private (noisy) SGD

Theorem

Assume that (i) $k \gg \sqrt{n}$*, (ii)* $m \leq \frac{n}{\sqrt{k}}$*, and (iii)* $1 \geq \varepsilon^2 \geq \frac{d}{m} \log \frac{1}{\delta}$*. Then DP-SGD achieves*

$$\mathbb{E}\left[f(\overline{\theta}_k) - f(\theta^\star)\right] \lesssim \frac{L_1 R^2}{k} + \frac{L_0 R}{\sqrt{km}}$$

Corollary (Bassily et al. [5, 6], Feldman et al. [10])

With appropriate choices for m *and* $k = \frac{n^2}{m^2}$*,*

$$\mathbb{E}\left[f(\widehat{\theta}_n) - f(\theta^\star)\right] \lesssim \frac{L_0 R \sqrt{d} \log \frac{1}{\delta}}{n\varepsilon}.$$

Amplification by shuffling

- individual data $x_1, \ldots, x_n$
- for differentially private Q, each reports

$$Z_i \sim Q(\cdot \mid x_i)$$

- permute $(Z_1, \ldots, Z_n)$ randomly

Intuition

- heuristic idea: suppose $x_1, \ldots, x_n \in \{0, 1\}$
- randomized response: for $q_\varepsilon = \frac{e^\varepsilon}{1+e^\varepsilon}$,

$$Z_i = \begin{cases} x_i & \text{w.p. } q_\varepsilon \\ 1 - x_i & \text{otherwise} \end{cases}$$

- key intuition: $\sum_{i=1}^n Z_i$ is sufficient statistic for $(Z_{\pi(i)})_{i=1}^n$
- using

$$\begin{aligned} \mathbb{E}[Z_i] &= q_\varepsilon x_i + (1 - q_\varepsilon)(1 - x_i) = (2q_\varepsilon - 1)x_i - (1 - q_\varepsilon) \\ \mathrm{Var}(Z_i) &= q_\varepsilon(1 - q_\varepsilon), \end{aligned}$$

approximately

$$\sum_{i=1}^n Z_i \dot{\sim} \mathsf{N}\left((2q_\varepsilon - 1)\mathbf{1}^T x - (1 - q_\varepsilon)n, nq_\varepsilon(1 - q_\varepsilon)\right)$$

Intuition (continued)

- for $x, x' \in \{0,1\}^n$ with $\|x - x'\|_1 \leq 1$,

$$\sum_{i=1}^{n} Z_i \dot{\sim} \mathsf{N}\left((2q_\varepsilon - 1)\mathbf{1}^T x - (1-q_\varepsilon)n, nq_\varepsilon(1-q_\varepsilon)\right)$$

$$\sum_{i=1}^{n} Z'_i \dot{\sim} \mathsf{N}\left((2q_\varepsilon - 1)\mathbf{1}^T x' - (1-q_\varepsilon)n, nq_\varepsilon(1-q_\varepsilon)\right)$$

- normals $\mathsf{N}(\mu, \sigma^2)$ and $\mathsf{N}(0, \sigma^2)$ are (ε, δ)-close if $\sigma^2 \gtrsim \mu^2 \frac{\log \frac{1}{\delta}}{\varepsilon^2}$
  - $\mu = (2q_\varepsilon - 1) = \frac{e^\varepsilon - 1}{e^\varepsilon + 1}$ and $\sigma^2 = n\frac{e^\varepsilon}{(1+e^\varepsilon)^2}$

final privacy

$$\varepsilon^2_{\mathsf{final}} = \frac{(e^\varepsilon - 1)^2 \log \frac{1}{\delta}}{ne^\varepsilon} = \frac{e^\varepsilon(1 - e^{-\varepsilon})^2 \log \frac{1}{\delta}}{n}.$$

A more formal statement

Theorem (Feldman et al. [11])

Let Z_i be ε_0-locally differentially private views of X_i and $\pi : [n] \to [n]$ a uniformly random permutation. Then $(Z_{\pi(1)}, \ldots, Z_{\pi(n)})$ is

$$\varepsilon = O(1)(1 - e^{-\varepsilon_0})\sqrt{\frac{e^{\varepsilon_0} \log \frac{1}{\delta}}{n}} \text{ differentially private}$$

Proof idea 1 (randomized response): releases as mixtures

$$Q(\cdot \mid x) = \frac{p}{2} Q(\cdot \mid 0) + \frac{p}{2} Q(\cdot \mid 1) + (1-p) \begin{bmatrix} x \\ 1-x \end{bmatrix}$$

Proof idea 2: simulate outputs

$$Y_i \in \{0,1,2\}, \quad Y_i \overset{\text{iid}}{\sim} \begin{cases} 0 & \text{w.p. } p/2 \\ 1 & \text{w.p. } p/2 \\ 2 & \text{w.p. } 1-p \end{cases}$$

and when $\pi(i) \neq 1$,

$$Z_i = \begin{cases} \mathsf{Ber}(q) & \text{if } Y_i = 0 \\ \mathsf{Ber}(1-q) & \text{if } Y_i = 1 \\ x_{\pi(i)} & \text{if } Y_i = 2 \end{cases}$$

Proof idea 3: control Binomial deviations

Bound difference between $(1 + N_0, N - N_0)$ and $(N_0, N - N_0 + 1)$ for

$$N \sim \mathsf{Bin}(n-1, p) \quad \text{and} \quad N_0 \mid N \sim \mathsf{Bin}(N, 1/2).$$

Other types of amplification and some references

- Binary case: Cheu et al. [7]
- RAPPOR system of Erlingsson et al. [8]
- Amplification by *iteration*, Feldman et al. [9] and Altschuler and Talwar [3]
- "application" in deep learning: moments-accountant, Abadi et al. [1]

A convergence analysis of stochastic gradient

$$\begin{array}{rl} \text{minimize} & f(x) \\ \text{subject to} & x \in X \end{array}$$

where f has L_1-Lipschitz gradient and stochastic gradients g satisfy

$$\mathrm{Var}(g \mid x) = \mathbb{E}[\|g - \nabla f(x)\|_2^2] \leq V^2$$

stochastic gradient iteration

$$x_{k+1} = \operatorname*{argmin}_{x \in X} \left\{ \nabla f(x_k)^T x + \frac{1}{2\eta} \|x - x_k\|_2^2 \right\}$$

Theorem

If $\eta^{-1} = L_1 + \gamma$, *then*

$$\sum_{i=1}^{k} \mathbb{E}[f(x_{i+1}) - f(x^\star)] \leq \frac{\|x_1 - x^\star\|^2}{2\eta} + \frac{k}{2\gamma} V^2.$$

Proof: building blocks

Idea: analyze one-step progress

Lemma (smoothness properties)

$$f(x_{k+1}) - f(x^\star) \leq \langle \nabla f(x_k), x_{k+1} - x^\star \rangle + \frac{L_1}{2} \|x_k - x_{k+1}\|^2 .$$

Lemma (intermediate minimizers)

for h convex, x minimizes h over X iff $\langle \nabla h(x), y - x \rangle \geq 0$, all $y \in X$

Proof: one-step progress

Lemma

for noise error $\xi_k = \nabla f(x_k) - g_k$,

$$f(x_{k+1}) - f(x^\star) \leq \frac{1}{2\eta}\left[\|x_k - x^\star\|^2 - \|x_{k+1} - x^\star\|^2 - \|x_k - x_{k+1}\|^2\right] + \langle \xi_k, x_{k+1} - x^\star \rangle + \frac{L_1}{2}\|x_k - x_{k+1}\|^2$$

Finalizing the analysis by telescoping

[1] M. Abadi, A. Chu, I. Goodfellow, B. McMahan, I. Mironov, K. Talwar, and L. Zhang. Deep learning with differential privacy. In *23rd ACM Conference on Computer and Communications Security (ACM CCS)*, pages 308–318, 2016.

[2] A. Agarwal and J. C. Duchi. Distributed delayed stochastic optimization. In *Advances in Neural Information Processing Systems 24*, 2011.

[3] J. Altschuler and K. Talwar. Privacy of noisy stochastic gradient descent: More iterations without more privacy loss. In *Advances in Neural Information Processing Systems 35*, 2022.

[4] B. Balle, G. Barthe, and M. Gaboardi. Privacy amplification by subsampling: Tight analyses via couplings and divergences. In *Advances in Neural Information Processing Systems 31*, pages 6277–6287, 2018.

[5] R. Bassily, A. Smith, and A. Thakurta. Private empirical risk minimization: Efficient algorithms and tight error bounds. In *55th Annual Symposium on Foundations of Computer Science*, pages 464–473, 2014.

[6] R. Bassily, V. Feldman, K. Talwar, and A. Thakurta. Private stochastic convex optimization with optimal rates. In *Advances in Neural Information Processing Systems 32*, 2019.

[7] A. Cheu, A. Smith, J. Ullman, D. Zeber, and M. Zhilyaev. Distributed differential privacy via shuffling. In *Advances in Cryptology: EUROCRYPT 2019*, 2019.

[8] U. Erlingsson, V. Pihur, and A. Korolova. RAPPOR: Randomized aggregatable privacy-preserving ordinal response. In *Proceedings of the 21st ACM Conference on Computer and Communications Security (CCS)*, 2014.

[9] V. Feldman, I. Mironov, K. Talwar, and A. Thakurta. Privacy amplification by iteration. In *59th Annual Symposium on Foundations of Computer Science*, 2018.

[10] V. Feldman, T. Koren, and K. Talwar. Private stochastic convex optimization: Optimal rates in linear time. In *Proceedings of the Fifty-Second Annual ACM Symposium on the Theory of Computing*, 2020.

[11] V. Feldman, A. McMillan, and K. Talwar. Hiding among the clones: A simple and nearly optimal analysis of privacy amplification by

shuffling. In *62nd Annual Symposium on Foundations of Computer Science*, 2021.

[12] J. Ullman. Rigorous approaches to data privacy. Lecture notes for CS7880, Northeastern University.