Differentially Private E-Values (Notes)

Daniel Csillag   Diego Mesquita

1. Background: Differential Privacy

- Intuition on DP: membership attacks, formalize datasets
- **Definition ((ε, δ) —DP).** A randomized algorithm $\mathcal{A}(\cdot)$ satisfies (ε, δ)-differential privacy if, for all (fixed) neighboring datasets D and D',

$$\mathbb{P}[\mathcal{A}(D) \in A] \leq e^{\varepsilon}\mathbb{P}[\mathcal{A}(D') \in A] + \delta, \qquad \text{for all sets } A.$$

- **Definition ((α, ε)-Rényi DP).** A randomized algorithm $\mathcal{A}(\cdot)$ satisfies (α, ε)-RDP for $\alpha > 1$ if, for all (fixed) neighboring datasets D and D',

$$D_\alpha(\mathcal{A}(D) \parallel \mathcal{A}(D')) \leq \varepsilon, \qquad \text{where} \qquad D_\alpha(P \parallel Q) := \frac{1}{\alpha - 1} \log \mathbb{E}_Q\left[\left(\frac{\mathrm{d}P}{\mathrm{d}Q}\right)^\alpha\right].$$

- Basic properties:
  - When $\alpha \to +\infty$ we recover $(\varepsilon, 0)$-DP
  - (α, ε)-RDP $\Longrightarrow$ for any $\delta > 0$, we are (ε', δ)-DP with $\varepsilon' = \varepsilon + \frac{\log(1/\delta)}{\alpha - 1}$.
- Fundamental properties:
  - **Composition:** If $\mathcal{A}_1(\cdot)$ is (α, ε_1)-RDP and $\mathcal{A}_2(\cdot)$ is (α, ε_2)-RDP,
    then $(\mathcal{A}_1(\cdot), \mathcal{A}_2(\cdot))$ is $(\alpha, \varepsilon_1 + \varepsilon_2)$-RDP.
  - **Post-processing:** If $\mathcal{A}(\cdot)$ is (α, ε)-RDP,
    then $f(\mathcal{A}(\cdot))$ is (α, ε)-RDP for any f.
  - **Corollary:** If f is bijective, then $\mathcal{A}(\cdot)$ is (α, ε)-RDP iff $f(\mathcal{A}(\cdot))$ is (α, ε)-RDP.
- **Basic example (Additive noise):** $\mathcal{A} : \text{Dataset} \to \mathbb{R}$
  Mechanism: $\mathcal{A}^{\text{DP}}(D) := \mathcal{A}(D) + \xi, \qquad \xi \sim \mathcal{N}(0, \sigma^2)$
  Goal: choose σ^2 to ensure (α, ε)-RDP

$$\begin{aligned} D_\alpha(\mathcal{A}^{\text{DP}}(D) \parallel \mathcal{A}^{\text{DP}}(D')) &= D_\alpha(\mathcal{N}(\mathcal{A}^{\text{DP}}(D), \sigma^2) \parallel \mathcal{N}(\mathcal{A}^{\text{DP}}(D'), \sigma^2)) \\ &= \frac{\alpha\left|\mathcal{A}^{\text{DP}}(D') - \mathcal{A}^{\text{DP}}(D)\right|^2}{2\sigma^2} \leq \frac{\alpha[\Delta(\mathcal{A})]^2}{2\sigma^2}, \\ &\qquad \text{where} \qquad \Delta(\mathcal{A}) = \sup_{D \text{ adj } D'} \left|\mathcal{A}^{\text{DP}}(D') - \mathcal{A}^{\text{DP}}(D)\right|. \end{aligned}$$

  Solution:

$$\sigma^2 = \frac{\alpha[\Delta(\mathcal{A})]^2}{2\varepsilon}.$$

- Also, Laplace mechanism exists.
- **Remark:** Classical DP paper: “Differentially Private X”

2. E-Values / Differentially Private E-Values

2.1. Hypothesis testing

- **Definition (E-Value).** Random variable $E(D)$ such that

$$E(D) \text{ is nonnegative} \qquad \text{and} \qquad H_0 \vdash \mathbb{E}[E(D)] \leq 1.$$

- **Goal:** $E(D) \rightsquigarrow E^{\mathrm{DP}}(D)$, which is (α, ε)-RDP.
- **Mechanism:**

$$E^{\mathrm{DP}}(D) := E(D) \cdot e^{-\xi} \qquad \text{with} \qquad \xi \perp\!\!\!\perp D.$$

- It's nonnegative
- Derive the required MGF bound: under H_0,

$$\mathbb{E}[E^{\mathrm{DP}}(D)] = \mathbb{E}[E(D)] \cdot \mathbb{E}\left[e^{-\xi}\right] \leq \mathbb{E}\left[e^{-\xi}\right] \overset{!}{\leq} 1.$$

- So we just need:
  - Some ξ ensuring RDP:

  $E^{\mathrm{DP}}(D)$ is (α, ε)-RDP $\iff \log E^{\mathrm{DP}}(D) = \log E^{\mathrm{DP}}(D) - \xi$ is (α, ε)-RDP
  - And satisfying the MGF bound

$$\mathbb{E}\left[e^{-\xi}\right] \leq 1.$$

- Log-sensitivity:

$$\Delta_{\log}(E) := \sup_{D \text{ adj } D'} |\log E(D) - \log E(D')|$$

- **Biased Gaussian mechanism:** deriving it:

We know that $\xi \sim \mathcal{N}\left(\mu, \alpha\left[\Delta_{\log}(E)\right]^2/2\varepsilon\right)$ ensures (α, ε)-RDP for any μ. So let's solve for μ:

$$\begin{aligned} \mathbb{E}\left[e^{-\xi}\right] = e^{-\mu+\alpha\left[\Delta_{\log}(E)\right]^2/4\varepsilon} &\leq 1 \\ \iff -\mu + \alpha\left[\Delta_{\log}(E)\right]^2/4\varepsilon &\leq 0 \\ \iff \alpha\left[\Delta_{\log}(E)\right]^2/4\varepsilon &\leq \mu \end{aligned}$$

So we get

$$\xi \sim \mathcal{N}\left(\frac{\alpha\left[\Delta_{\log}(E)\right]^2}{4\varepsilon}, \frac{\alpha\left[\Delta_{\log}(E)\right]^2}{2\varepsilon}\right).$$

- **Biased Laplace mechanism:** For any $\alpha > 1$ and $\varepsilon > 0$, let

$$\xi \sim \mathrm{Lap}\left(-\log\left(1 - b_{\alpha,\varepsilon}^2\right), b_{\alpha,\varepsilon}\right),$$

where

$$\begin{aligned} b_{\alpha,\varepsilon} &= 1/\boldsymbol{h}_\alpha^{-1}\left((2\alpha-1)e^{(\alpha-1)\varepsilon}\right), \\ \boldsymbol{h}_\alpha(t) &= \alpha e^{(\alpha-1)\Delta_{\log}(E)t} + (\alpha-1)e^{-\alpha\Delta_{\log}(E)t}, \quad t \geq 0, \end{aligned}$$

if $\boldsymbol{b_{\alpha,\varepsilon} < 1}$.

  - Laplace MGF:

$$\xi \sim \mathrm{Lap}(\mu, b) \qquad \Longrightarrow \qquad \mathbb{E}\left[e^{t\xi}\right] = \frac{e^{\mu t}}{1 - b^2 t^2} \text{ for } |t| \leq 1/b.$$

- **Power:**

$$\mathbb{E}\left[\frac{1}{n}\log E^{\mathrm{DP}}(D)\right] = \mathbb{E}\left[\frac{1}{n}\log E(D)\right] - \frac{\mathbb{E}[\xi]}{n}$$

- Properties:
  - **Optional continuation:** $E_1^{\mathrm{DP}}(D_1)$ and $E_2^{\mathrm{DP}}(D_2)$ are (α, ε)-RDP $\Longrightarrow$ $\left(E_1^{\mathrm{DP}}(D_1), E_1^{\mathrm{DP}}(D_1) \cdot E_2^{\mathrm{DP}}(D_2)\right)$ is (α, ε)-RDP.
  - **E-to-p conversion:** $E^{\mathrm{DP}}(D)$ is (α, ε)-RDP $\Longrightarrow 1/E^{\mathrm{DP}}(D)$ also is
  - Averaging:
    - Independent averaging: $E_1^{\mathrm{DP}}(D_1)$ and $E_2^{\mathrm{DP}}(D_2)$ are (α, ε)-RDP $\Longrightarrow \eta E_1^{\mathrm{DP}}(D_1) + (1-\eta)E_2^{\mathrm{DP}}(D_2)$ is (α, ε)-RDP for any $\eta \in [0, 1]$
    - Dependent averaging: $E_1^{\mathrm{DP}}(D)$ and $E_2^{\mathrm{DP}}(D)$ are (α, ε)-RDP $\Longrightarrow \eta E_1^{\mathrm{DP}}(D) + (1-\eta)E_2^{\mathrm{DP}}(D)$ is $(\alpha, 2\varepsilon)$-RDP for any $\eta \in [0, 1]$

2.2. Algorithms atop e-values and confidence intervals

- Core result:
  - **Assumption.** If $E_1, ..., E_k$ are e-values for nulls $H_0^{(1)}, ..., H_0^{(k)}$, then $\mathcal{A}(E_1, ..., E_k)$ is *valid.*
  - **Theorem.** Under the assumption, $\mathcal{A}\left(E_1^{\mathrm{DP}}, ..., E_k^{\mathrm{DP}}\right)$ is *valid.*

- For confidence intervals, problem: finite k!

- Solution: reduce to finite k.

  If each $\log E_\theta(D)$ is locally Lipschitz wrt θ, partition $[\theta_{\mathrm{inf}}, \theta_{\mathrm{sup}}]$ into uniform bins, and within each bin:

$$\begin{aligned}
\log E_{\theta'}(D) &\geq \log \underbrace{E_{\theta_j}}_{\theta_j \text{ is the midpoint}}(D) - L_j|\theta' - \theta_j| \\
\Longrightarrow E_{\theta'}(D) &\geq E_{\theta_j}(D) \cdot e^{-L_j|\theta' - \theta_j|} \\
&\geq E_{\theta_j}(D) \cdot e^{-L_j \frac{a_j - a_{j-1}}{2}} =: \tilde{E}_{\theta_j}(D). \\
&\geq E_{\theta_j}(D) \cdot e^{-L_j(a_j - a_{j-1})} =: \tilde{E}_{\theta_j}(D).
\end{aligned}$$

  This is an e-value for all θ in the bin.

3. Concrete e-values

- **Mean e-value** with Cover's universal portfolios; $[\lambda_{\mathrm{inf}}, \lambda_{\mathrm{sup}}] \subsetneq \left(-1/(1-\theta), \frac{1}{\theta}\right)$.
  - Invariance to observations, this is very useful
  - **Proposition:**

$$\Delta_{\log}(E_\theta) \leq \max\left\{\log\left(1 + \max\{\lambda_{\mathrm{sup}}(1-\theta), -\lambda_{\mathrm{inf}}\theta\}\right), -\log\left(1 + \min\{\lambda_{\mathrm{inf}}(1-\theta), -\lambda_{\mathrm{sup}}\theta\}\right)\right\}.$$

  - **Proposition:**

$$\|\theta \mapsto \log E_\theta(D)\|_{\mathrm{Lip}} \leq \max\left\{\left|\frac{\lambda_{\mathrm{sup}}}{1 - \lambda_{\mathrm{sup}}\theta_{\mathrm{sup}}}\right|, \left|\frac{\lambda_{\mathrm{inf}}}{1 - \lambda_{\mathrm{inf}}(1 - \theta_{\mathrm{inf}})}\right|\right\}, \qquad \text{for } \theta \in [\theta_{\mathrm{inf}}, \theta_{\mathrm{sup}}]$$

- **Conformal e-prediction**

$$C_\alpha(x) = \{y \in \mathcal{Y} : E^{\mathrm{exch}}(D; s(x, y)) < 1/\alpha\}$$

$$E^{\mathrm{exch}}(D; S^{\mathrm{test}}) = \frac{(n+1)S^{\mathrm{test}}}{\sum_{i=1}^{n} s(X_i, Y_i) + S^{\mathrm{test}}}$$

- **Proposition:** If the image of s is in $[a, b]$, then

$$\Delta_{\log}\left(E^{\mathrm{exch}}\right) \leq 2 \cdot \frac{b/a}{n+1}.$$

To make finite: quantize s so that its image is finite