Differential Privacy as a Mutual Information Constraint

Paul Cuff and Lanqing Yu

Database Privacy

- Let $X_1, X_2, \ldots, X_n$ be entries in a database
  - E.g. X_i is personal information about person i
- Let Y be the response to a query
- The job of the information provider is to answer queries and protect individual privacy

Design $P(y|x)$

Differential Privacy

- ε-DP:
  - Let x and x′ differ in only one entry (i.e. $x_i = x'_i$ for all but one i)
  - $p(y|x) \leq e^{\varepsilon} p(y|x')$
- Why x and x′ differ in only one spot?
  - Convince someone to put their data in your database
- Why multiplicative constraint?
  - Posterior update is small

Dwork, 2006

A Common Technique

- Add Laplacean noise

Weaker DP

- (ε,δ)-DP:
  - Let x and x′ differ in only one entry
  - $P(Y \in A|x) \leq e^{\varepsilon} P(Y \in A|x') + \delta$
- Additive Gaussian noise often provides privacy

Dwork, 2006

Mutual Information Differential Privacy

ε-MI-DP:

$$\max_{i, P_{X^n}} I(X_i; Y | X^{i-1}, X_{i+1}^n) < \epsilon$$

Claim

ε-DP > MI-DP > (ε,δ)-DP

Furthermore, if input or output alphabet is finite,

MI-DP = (ε,δ)-DP

similar to semantic security proof, Bellare, Tessaro, Vardy, 2012

Privacy Ordering

- α-DP > β-DP if for all $\beta>0$ there exists α such that α-DP $\Rightarrow$ β-DP.

Subadditivity of DP

- Multiple queries:
  - If k queries $Y_{q1}, Y_{q2}, \ldots, Y_{qk}$ each have differential privacy ε and are conditionally independent, the combined they have kε privacy.

Simple MI-DP Proof:

$$
\begin{aligned}
I(X;Y_1,Y_2) &= I(X;Y_1) + I(X;Y_2|Y_1) \\
&\leq I(X;Y_1) + I(X;Y_2)
\end{aligned}
$$

For clarity, conditioned database variables are omitted.

Common complaint

- Differentially privacy doesnt not mean that you cant learn about X_i.
    - Consider a database with correlated entries.

Simple MI-DP Explanation:

$$I(X_i;Y) \quad \nleq \quad I(X_i;Y|X^{i-1},X_{i+1}^n)$$

Precise Bounds

(ε,δ)-Closeness

$$P \overset{(\epsilon,\delta)}{\approx} Q$$

if

$$P(A) \le e^{\epsilon} Q(A) + \delta, \quad \forall A \in \mathcal{F},$$
$$Q(A) \le e^{\epsilon} P(A) + \delta, \quad \forall A \in \mathcal{F}.$$

Special Cases

$$P \overset{(\epsilon,0)}{\approx} Q \iff \left|\ln \frac{dP}{dQ}(a)\right| \leq \epsilon \quad \forall a \in \Omega.$$

$$P \overset{(0,\delta)}{\approx} Q \iff \|P - Q\|_{TV} \leq \delta.$$

Relation to Detection Theory

$$P \overset{(\epsilon,\delta)}{\approx} Q$$

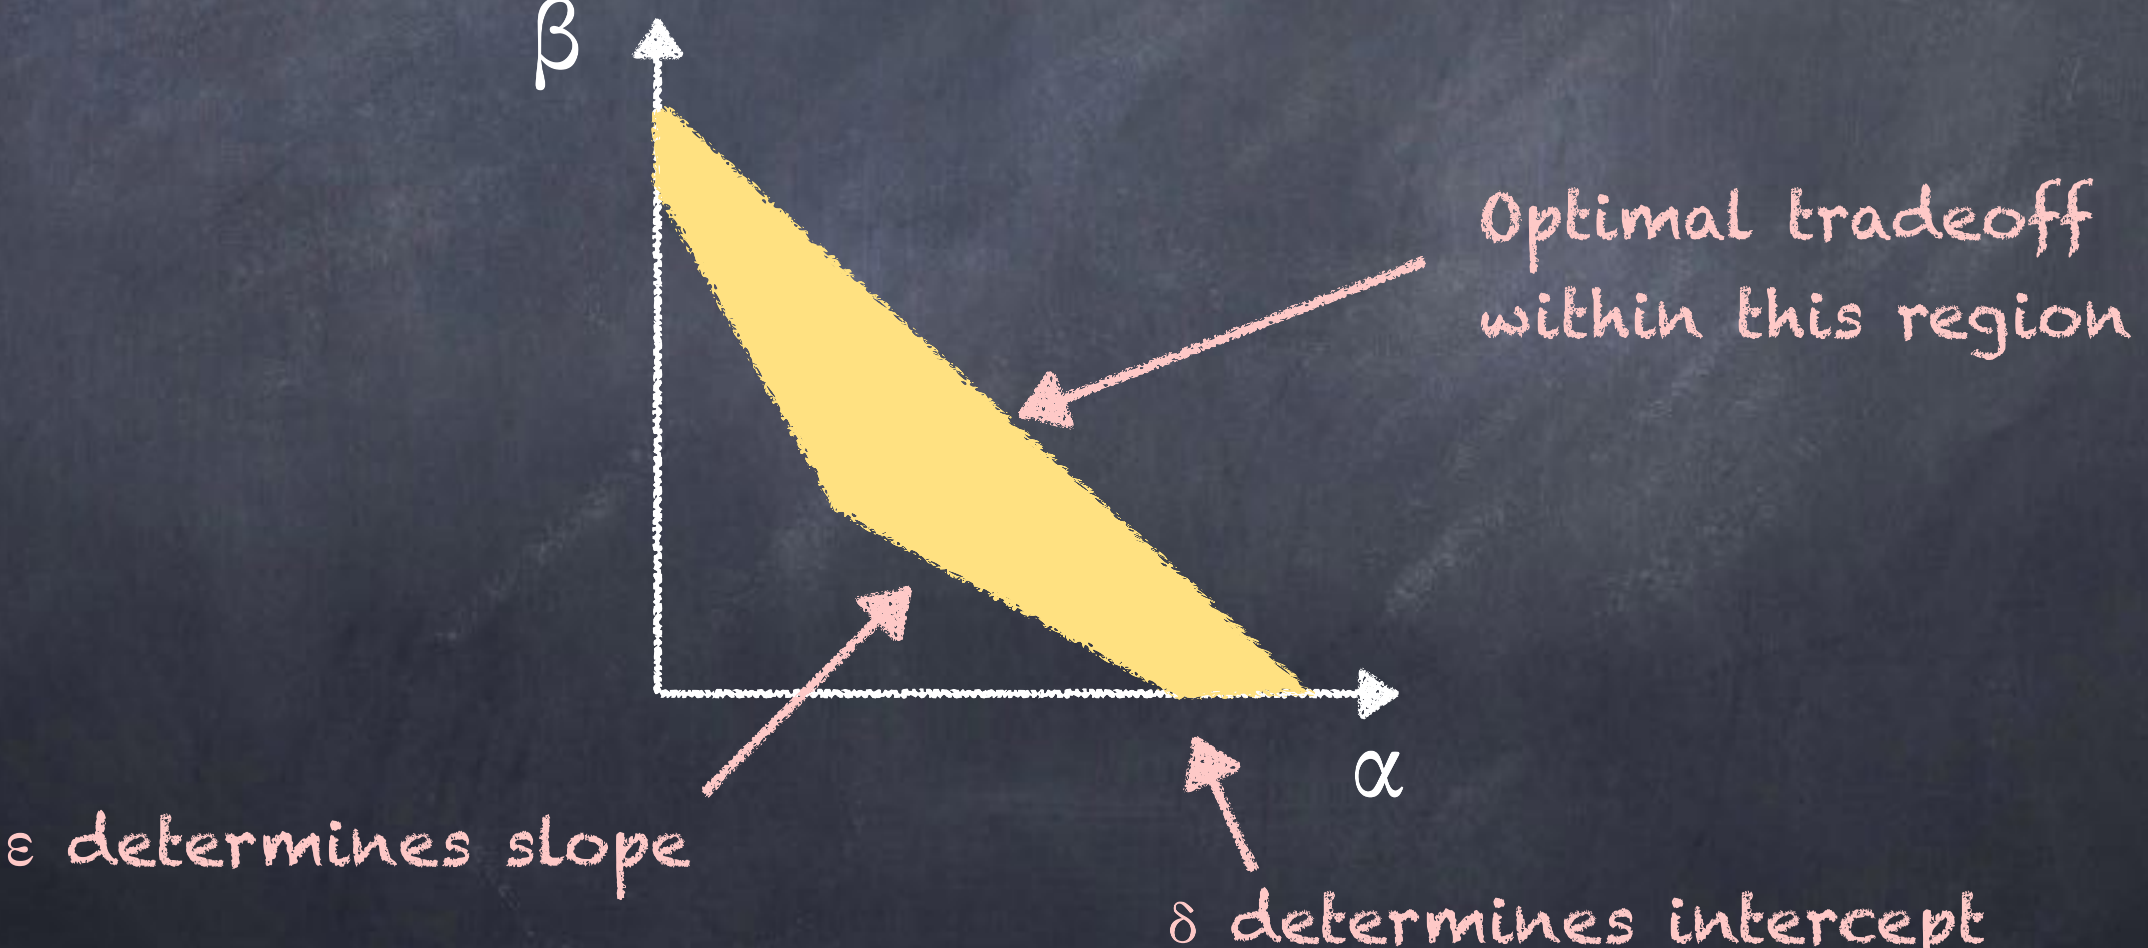

Kairouz, Oh, Viswanath, 2015

Tightest (ε,δ) Conversion

$$P \overset{(\epsilon,\delta)}{\approx} Q \implies P \overset{\left(\epsilon',\delta'\right)}{\approx} Q.$$

for

$$\epsilon' \leq \epsilon$$

$$\delta' = 1 - \frac{\left(e^{\epsilon'} + 1\right)(1 - \delta)}{e^{\epsilon} + 1}$$

Simple Claim

$$P \overset{(\epsilon,0)}{\approx} Q \quad \Longrightarrow \quad \begin{aligned} D(P\|Q) &\le \epsilon \text{ nats}, \\ D(Q\|P) &\le \epsilon \text{ nats}. \end{aligned}$$

Tightest Claim

$$P \overset{(\epsilon,0)}{\approx} Q \quad \Longrightarrow \quad \begin{aligned} D(P\|Q) &\le \epsilon \frac{(e^{\epsilon}-1)(1-e^{-\epsilon})}{(e^{\epsilon}-1)+(1-e^{-\epsilon})} \text{ nats}, \\ D(Q\|P) &\le \epsilon \frac{(e^{\epsilon}-1)(1-e^{-\epsilon})}{(e^{\epsilon}-1)+(1-e^{-\epsilon})} \text{ nats}. \end{aligned}$$

Pinsker

$$D(P\|Q) \le \epsilon \text{ nats} \quad \Longrightarrow \quad P \overset{\left(0,\sqrt{\epsilon/2}\right)}{\approx} Q.$$

Relative entropy to Mutual Information

If

$$D\left(P_{Y|X=x_1} \| P_{Y|X=x_2}\right) \leq \epsilon \quad \forall x_1, x_2 \in \mathcal{X}$$

then

$$I(X;Y) \leq \epsilon$$

Hint: Radius of information ball

Mutual Information to Total Variation

$$\max_{P_X} I(X;Y) \leq \epsilon \quad \Longrightarrow \quad \begin{array}{c} \|P_{Y|X=x_1} - P_{Y|X=x_2}\|_{TV} \leq \delta' \\ \forall x_1, x_2 \in \mathcal{X} \end{array}$$

$$\begin{aligned} \delta' &= 1 - 2h^{-1}(\ln 2 - \epsilon) \\ &\leq \sqrt{2\epsilon} \end{aligned}$$

Tightest bound, achieved with binary channel

Finite Alphabet

$$\left\|P_{Y|X=x_1} - P_{Y|X=x_2}\right\|_{TV} \le \delta \quad \forall x_1, x_2 \in \mathcal{X} \implies I(X;Y) \le \epsilon'$$

$$\epsilon' = 2h(\delta) + 2\delta \ln\left(\min\left\{|\mathcal{Y}|, \max_i |\mathcal{X}_i| + 1\right\}\right)$$

harder step

Continuity of entropy

Continuity of conditional entropy

inspired by Alicki and Fannes, 2004

Observation

$$\max_{P_{X^n}} I(X_i;Y|X^{i-1},X_{i+1}^n) = \max_{\prod_{t=1}^n P_{X_t}} I(X_i;Y) \quad \forall i$$

Either could be used for definition of MI-DP

Pointed out by Thomas Steinke