Measuring privacy and utility of local privacy protocols

Milan Lopuhaä-Zwakenberg[1]

[1]Security group, Technische Universiteit Eindhoven

May 13, 2019

Setup

Scenario:

- Everyone has private data (e.g. “are you an alcoholic” yes/no);
- Some aggregator (researcher, tech company) wants statistics about the population as a whole;
- Individuals do not trust aggregator with their private data.

Setup

Solution:

- Give random answers! (but let probabilities depend on private data)
- E.g. Yes/no-question: $\mathbb{P}(\text{tell truth}) = 3/4$, $\mathbb{P}(\text{lie}) = 1/4$.
- Aggregator cannot be certain about any individual's private data, but can reconstruct original statistics.

Abstract formulation

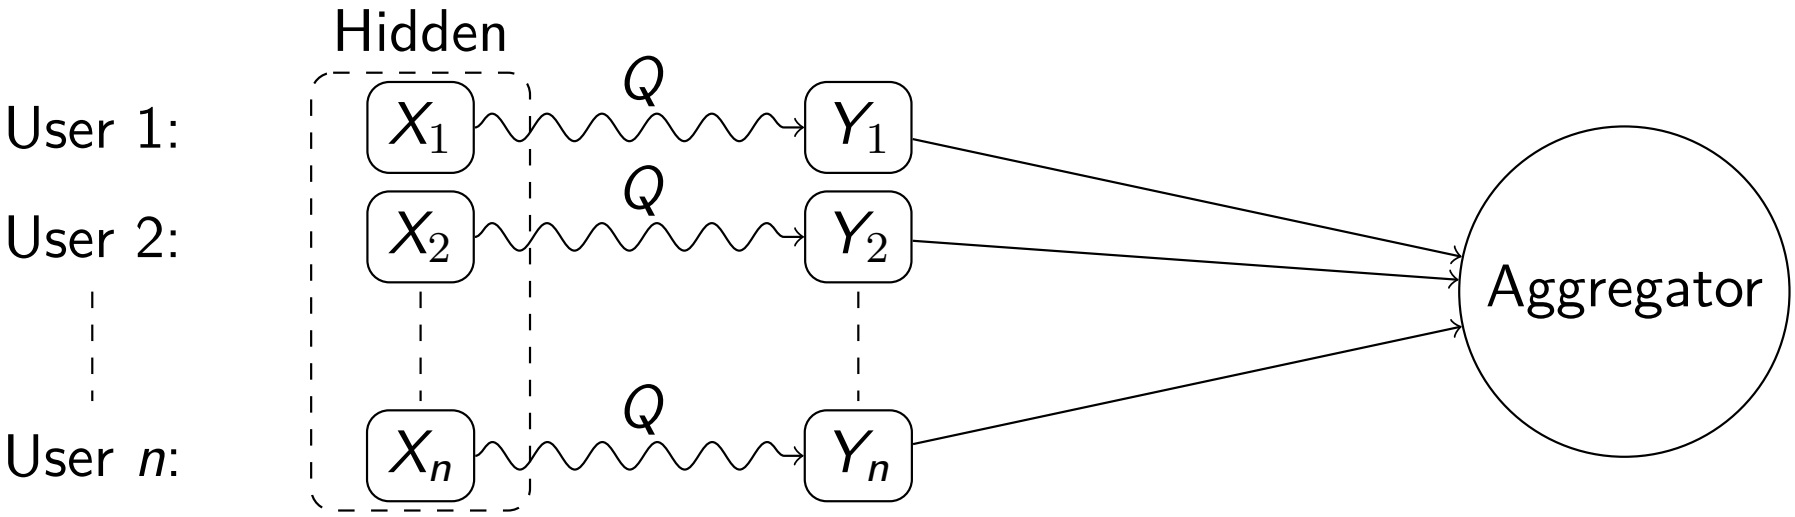

- private data $X_1, \cdots, X_n$ from (finite) set $\mathcal{A} = \{1, \cdots, a\}$;
- random function Q: $\mathcal{A} \rightsquigarrow \mathcal{B} = \{1, \cdots, b\}$;

$$Q = \begin{pmatrix} Q_{1|1} & Q_{1|2} & \cdots & Q_{1|a} \\ Q_{2|1} & Q_{2|2} & \cdots & Q_{2|a} \\ \vdots & \vdots & \ddots & \vdots \\ Q_{b|1} & Q_{b|2} & \cdots & Q_{b|a} \end{pmatrix} \text{ with } Q_{y|x} = \mathbb{P}(Q(X_i) = y | X = x).$$

Abstract formulation

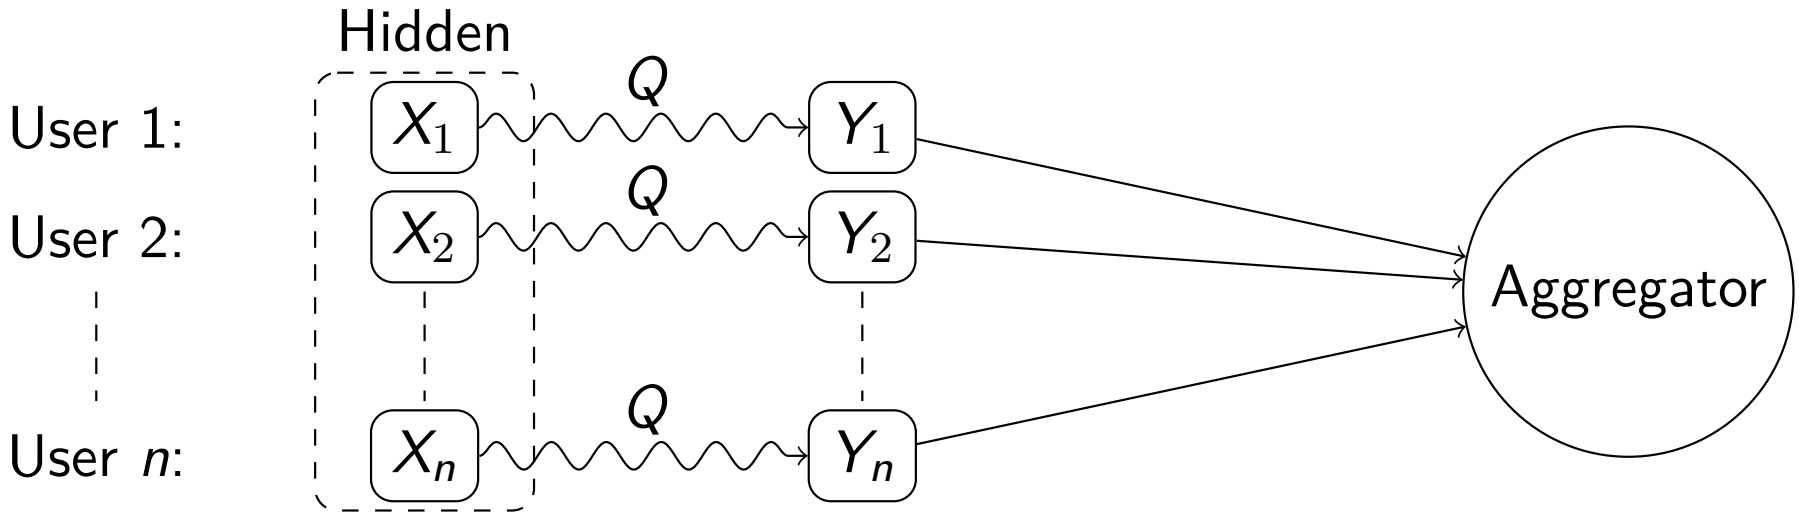

Goals:

- Aggregator: learn about $f_x = \frac{1}{n}\#\{i : X_i = x\}$ for all x.
- Users: prevent the aggregator from learning about their individual X_i.

Intuition:

- The more 'random' Q is, the more privacy for the users, but the less utility for the aggregator.

Our research

"The more 'random' Q is, the more privacy for the users, but the less utility for the aggregator."

Question (Cool people)

How can we make privacy protocols $(Q, \mathcal{B})$ that offer sufficient privacy to the users, and maximise utility for the aggregator?

Question (Me)

How do we define *privacy* and *utility*?

Privacy

Definition

Let $\varepsilon \geq 0$. Q satisfies ε-*Local Differential Privacy* if for all $x, x' \in \mathcal{A}$, $y \in \mathcal{B}$:

$$Q_{y|x} \leq \mathrm{e}^{\varepsilon} Q_{y|x'}.$$

We define $\mathrm{LDP}(Q) := \min\{\varepsilon : Q \text{ satisfies } \varepsilon\text{-LDP}\}$.

Problems:

- Very strict: even for $y \in \mathcal{B}$ with very low probability we have requirements.
- Does not pick up on many things we would like to consider as privacy.
- With what ε would a user be satisfied?

Utility

For $y \in \mathcal{B}$, let $s_y := \frac{1}{n}\#\{i : Y_i = y\}$. For $x \in \mathcal{A}$, take an estimator $\hat{f}_x : \mathbb{R}^{\mathcal{B}} \to \mathbb{R}$, and define utility as

$$\mathbb{E}\left[\sum_{x \in \mathcal{A}} (\hat{f}_x((s_y)_{y \in \mathcal{B}}) - f_x)^2 \middle| X_1, \cdots, X_n\right]. \tag{1}$$

Problems:

- Clearly depends on choices of $\hat{f}_x$ and the values of the X_i.
- In our yes/no example, $\hat{f}_{\text{yes}} = 2\#\{i : Y_i = \text{yes}\} - \frac{1}{2}$. How to deal with negative values?

Information Theory

- A probability event: *information content* of A is $\mathrm{i}(A) = -\log(\mathbb{P}(A))$.
- X discrete random variable on $\mathcal{A}$: *entropy* of X is

$$\mathrm{H}(X) = \mathbb{E}_x\left[\mathrm{i}(X = x)\right] = -\sum_{x \in \mathcal{A}} \mathrm{p}_x \log(\mathrm{p}_x);$$

  measure of the information stored in X or uncertainty about the value of X.
- Y another random variable: *conditional entropy* of X given Y is

$$\mathrm{H}(X|Y) = \mathbb{E}_y\left[\mathrm{H}(X)|Y = y\right] = -\sum_{x,y} \mathrm{p}_{x,y} \log(\mathrm{p}_{x|y}).$$

  measure of the uncertainty about X left after learning Y.
- One has $\mathrm{H}(X|Y) \leq \mathrm{H}(X)$ with equality iff X and Y are independent.

Information Theory

- Mutual information: $\mathrm{I}(X;Y) = \mathrm{H}(X) - \mathrm{H}(X|Y)$, amount of information learned about X by knowing Y (or vice versa).
- Conditional mutual information: $\mathrm{I}(X;Y|Z) = \mathrm{H}(X|Z) - \mathrm{H}(X|Y,Z)$.

Reframing

To describe privacy and utility in terms of information theory, we need to describe everything in a probabilistic framework.

- All the X_i are random variables, drawn independently from the same p.d. $\bar{P} = (P_1, \cdots, P_a)$ on $\mathcal{A}$.
- We consider $\bar{P}$ as a continuous random variable on

$$\mathcal{P}_{\mathcal{A}} = \left\{ \bar{p} \in \mathbb{R}^{\mathcal{A}}_{\geq 0} : \sum_{x \in \mathcal{A}} p_x = 1 \right\};$$

  its distribution Δ reflects prior knowledge about $\bar{P}$.
- The aggregator wants to learn $\bar{P}$.

Definitions

We can decompose

$$\underbrace{\mathrm{H}(\vec{X})}_{\text{total info in } \vec{X}} = \mathrm{I}(\vec{X};\bar{P}) + \mathrm{H}(\vec{X}|\bar{P}) = \underbrace{\mathrm{I}(\vec{X};\bar{P})}_{\text{nonprivate info}} + \sum_{i=1}^{n} \underbrace{\mathrm{H}(X_i|\bar{P})}_{i\text{'s private info}},$$

$$\underbrace{\mathrm{I}(\vec{X};\vec{Y})}_{\substack{\text{total info in } \vec{X} \\ \text{available} \\ \text{to aggregator}}} = \mathrm{I}(\vec{Y};\bar{P}) + \mathrm{H}(\vec{X}|\vec{Y},\bar{P}) = \underbrace{\mathrm{I}(\vec{Y};\bar{P})}_{\substack{\text{nonprivate info} \\ \text{available} \\ \text{to aggregator}}} + \sum_{i=1}^{n} \underbrace{\mathrm{I}(X_i;Y_i|\bar{P})}_{\substack{i\text{'s private info} \\ \text{available} \\ \text{to aggregator}}}.$$

Definitions

Utility:

$$\mathrm{Uti}_{n,\Delta}(Q) = \frac{\mathrm{I}(\vec{Y};\bar{P})}{\mathrm{I}(\vec{X};\bar{P})} = \begin{array}{c}\text{part of nonprivate info}\\ \text{available to aggregator.}\end{array}$$

Privacy:

$$\mathrm{Priv}_{\Delta}(Q) = \frac{\mathrm{H}(X_i|\,Y_i,\bar{P})}{\mathrm{H}(X_i|\bar{P})} = 1 - \frac{\mathrm{I}(X_i;\,Y_i|\bar{P})}{\mathrm{H}(X_i|\bar{P})} = \begin{array}{c}\text{part of private info}\\ \text{hidden from aggregator.}\end{array}$$

Note: both depend on distribution Δ of $\bar{P}$.

Privacy

$$\mathrm{Priv}_{\Delta}(Q) = \frac{\mathrm{H}(X_i|Y_i,\bar{P})}{\mathrm{H}(X_i|\bar{P})} = 1 - \frac{\mathrm{I}(X_i;Y_i|\bar{P})}{\mathrm{H}(X_i|\bar{P})} = \begin{array}{c}\text{part of private info}\\ \text{hidden from aggregator.}\end{array}$$

LDP problems:

- Very strict: even for $y \in \mathcal{B}$ with very low probability we have requirements. – *Solved by 'averaging' over y*
- Does not pick up on many things we would like to consider as privacy. – *Solved: any information protection is measured*
- With what ε would a user be satisfied? – *Solved: gives a %-age of hidden information*

Utility

$$\mathrm{Uti}_{n,\Delta}(Q) = \frac{\mathrm{I}(\vec{Y};\bar{P})}{\mathrm{I}(\vec{X};\bar{P})} = \frac{\text{part of nonprivate info}}{\text{available to aggregator.}}$$

Estimator problems:

- Clearly depends on choices of $\hat{f}_x$ and the values of the X_i. *Solved: no estimator involved, average over X_i*
- In our example, $\hat{f}_{\text{yes}} = 2 \cdot s_{\text{yes}} - \frac{1}{2}$. How to deal with negative values? *Solved: no estimator involved, instead one can compute posterior distributions*

Utility theorems

Reminder: Q can be seen as a $(b \times a)$-matrix.

Theorem

1. $\lim_{n\to\infty} \mathrm{Uti}_{n,\Delta}(Q) = \frac{\mathrm{rk}(Q)-1}{a-1}$.
2. *If* $\mathrm{rk}(Q) = a$, *then* $\exists$ *computable constants* b_Δ, $\mathrm{a}_\Delta(Q)$ *s.t.*

$$\mathrm{Uti}_{n,\Delta}(Q) \approx 1 - \frac{\mathrm{b}_\Delta - \mathrm{a}_\Delta(Q)}{\log n}.$$

3. *For large n:*

$$\mathrm{I}(\underbrace{\vec{Y}}_{n}; \bar{P}) \approx \mathrm{H}(\bar{P}\langle \tfrac{1}{2}\log n + \mathrm{a}_\Delta(Q)\rangle),$$

where $\bar{P}\langle\eta\rangle$ *is the discretisation of* $\bar{P}$ *in* η *digits.*

So goal: find Q that maximises $\mathrm{a}_\Delta(Q)$ while satisfying privacy constraints.

Privacy theorems

Theorem

Let Q be a privacy protocol. Then

$$\inf_{\Delta} \inf_{y} \frac{\mathrm{H}(X_i|\, Y_i = y, \bar{P})}{\mathrm{H}(X_i|\bar{P})} = e^{-\,\mathrm{LDP}(Q)}, \tag{2}$$

$$\mathrm{Priv}_{\Delta}(Q) = \frac{\mathrm{H}(X_i|\, Y_i, \bar{P})}{\mathrm{H}(X_i|\bar{P})} \geq e^{-\,\mathrm{LDP}(Q)}. \tag{3}$$

Hence LDP measures *worst case privacy*, Priv measures *average privacy*.

Comparing protocols

We can use our metrics to compare privacy protocols (GRR vs UE):

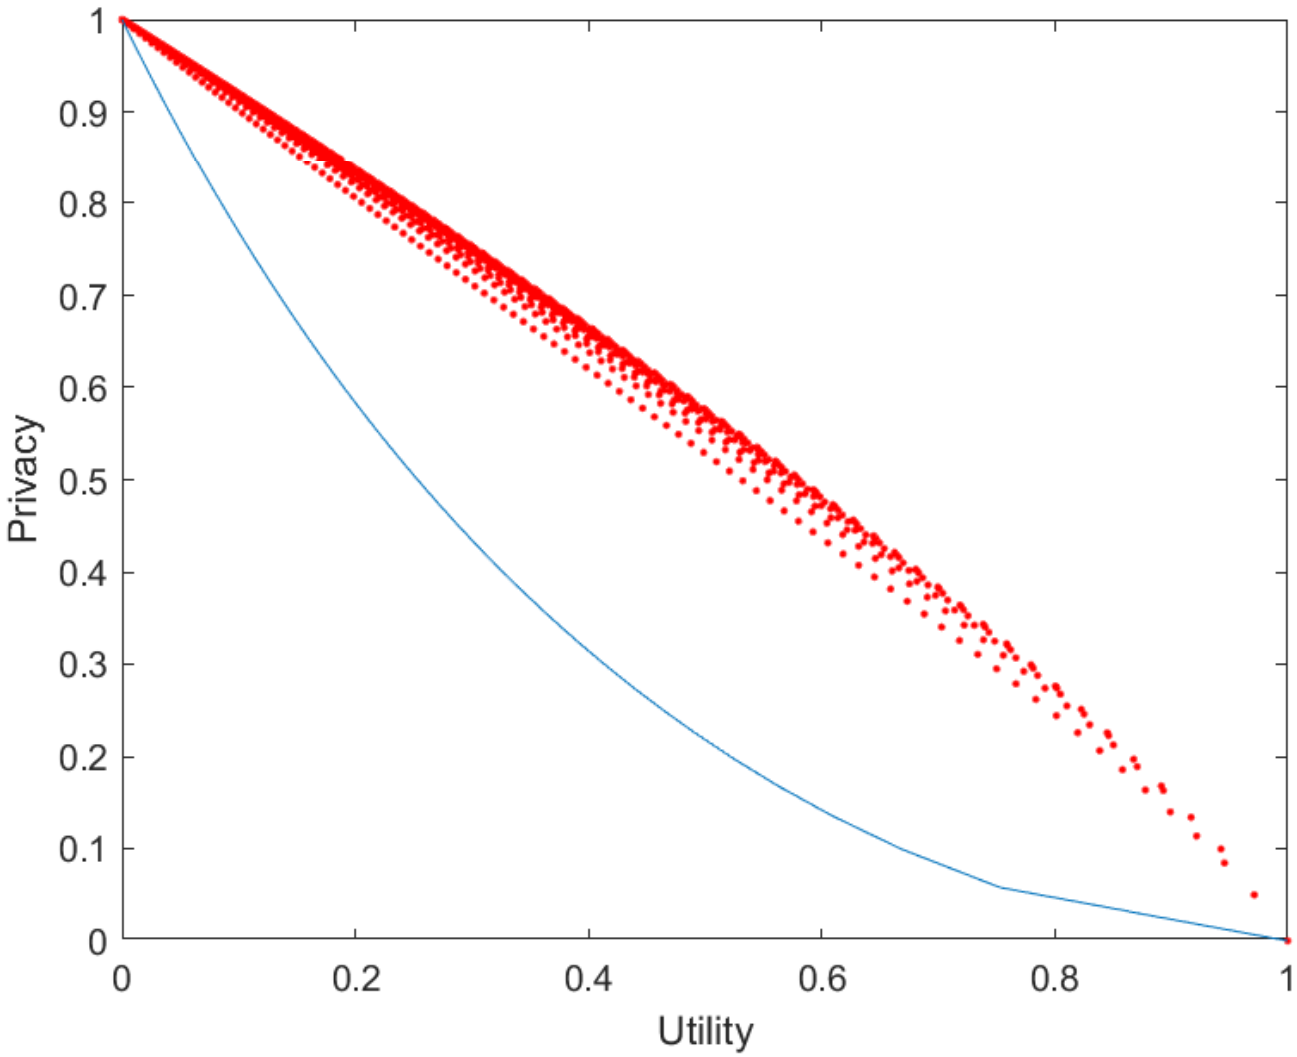

Work in progress

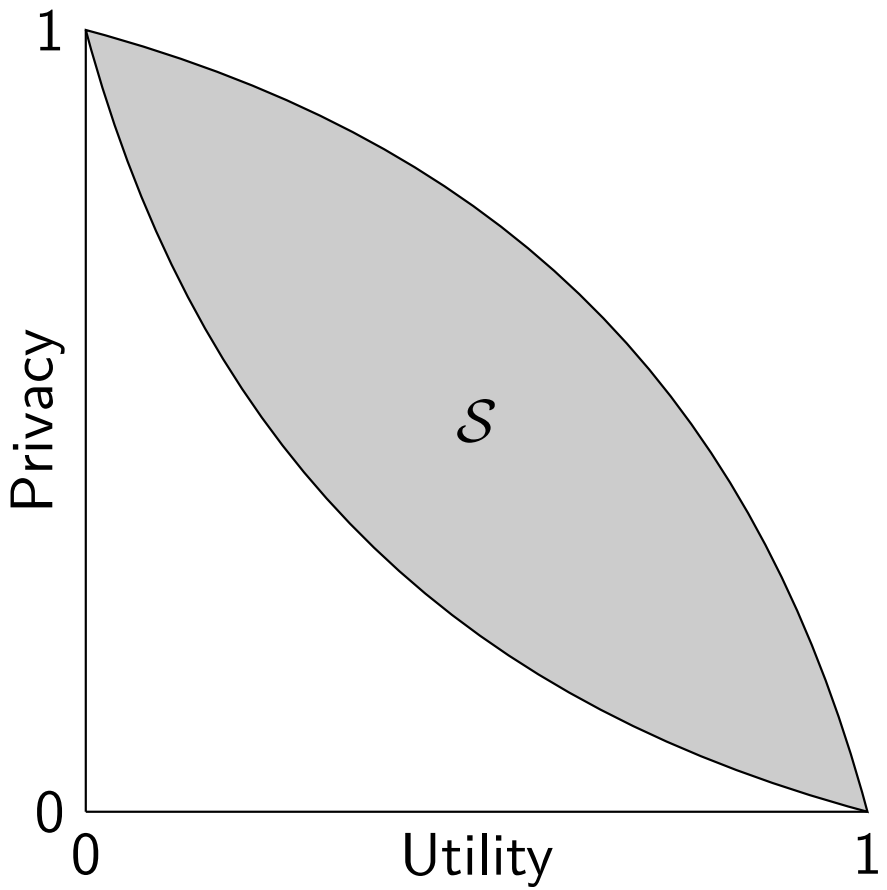

Can we describe $\mathcal{S}$ or the protocols on the boundary?

Conclusion

- Information theory gives intuitive metrics for privacy and utility;
- These metrics don't have the problems old metrics have;
- We can give 'meaning' to LDP;
- To do: formalise privacy-utility tradeoff.