1 Differential Privacy

1.1 Setting: Collecting and Providing Statistical Data

- Census bureau: Income distributions, "How many people earn > \$100,000?". Hospitals: Statistics about medical conditions, "How many smokers among pancreatic-cancer patients?", etc.
- Problem: Gender, age, weight, ethnicity, and marital status (for example) may be sufficient for identification among 1000 patients. Illustrative example: "AOL search data scandal"

1.2 An Utopian Goal [3]

- Ideally: Cannot learn anything about an individual that could not be learned without access to statistical database [1]. Similar to *semantic security*: Nothing can be learned about a plaintext from the ciphertext that could not be learned without seeing the ciphertext [6].
- Impossible in this generality (while semantic security is possible). Example: *Statistical information* on average income. *Auxiliary information*: 20% higher than average. Note: Additional information gives solution regardless of person is in database.

1.3 The Differential-Privacy Approach [4]

- Other approach to privacy: (Non-)Participating in a statistical database does not substantially affect the outcome of any analysis.
- Model: A *database* $\boldsymbol{d}$ is a string $d_1, \ldots, d_n$ of length n over some set D. Each d_i is called a *row* in $\boldsymbol{d}$. Two databases are *neighbors* if they coincide in $(n-1)$ rows. A *query* is a function $f : D^n \to \mathcal{R}$. For now, assume $\mathcal{R} \subset \mathbb{R}$ bounded. A *privacy mechanism* K_f adds noise to the true answer $f(d)$ to produce the *response* $K_f(d) = f(d) + \Delta$, where Δ is a random variable.
- **Definition**: K_f gives ε-differential privacy if for all neighboring $\boldsymbol{d}, \boldsymbol{d'} \in D^n$ and all $S \subseteq \mathcal{R}$, $\Pr[K_f(\boldsymbol{d}) \in S] \leq \exp(\varepsilon) \cdot \Pr[K_f(\boldsymbol{d'}) \in S]$.
- **Definition**: The *sensitivity* of a query $f : D^n \to \mathcal{R}$ is $\Delta f := \max_{\text{neighbors } \boldsymbol{d},\boldsymbol{d'}} |f(\boldsymbol{d}) - f(\boldsymbol{d'})|$.
- Typical privacy mechanism: Let $K_f = f(\boldsymbol{d}) + \Delta$, where $\Delta \sim \mathrm{Lap}(0, b)$ (Laplacian distribution with mean 0 and variance $2b^2$) and $b = \Delta f / \varepsilon$, i.e., with densitiy

  $$x \mapsto \frac{1}{2b} \exp\left(\frac{-|x|}{b}\right)$$

  and cumulative distribution function $\Pr[\Delta \leq x] = \frac{1}{2}\left[1 + \mathrm{sgn}(x)\left(1 + \exp\left(\frac{-|x|}{b}\right)\right)\right]$.
- This gives ε-differential privacy:

  $$\begin{aligned}
  \Pr[K_f(\boldsymbol{d}) \in S] &= \int_S \frac{\varepsilon}{2\Delta f} \exp\left(\varepsilon \frac{-|f(\boldsymbol{d}) - r|}{\Delta f}\right) d\lambda(r) \\
  &\leq \int_S \frac{\varepsilon}{2\Delta f} \exp\left(\varepsilon \frac{\Delta f - |f(\boldsymbol{d'}) - r|}{\Delta f}\right) d\lambda(r) \\
  &\qquad \left[-\Delta f - |f(\boldsymbol{d}) - r| \leq -|f(\boldsymbol{d'}) - f(\boldsymbol{d})| - |f(\boldsymbol{d}) - r| \leq -|f(\boldsymbol{d'}) - r|\right] \\
  &= \exp(\varepsilon) \cdot \Pr[K_f(\boldsymbol{d'}) \in S]
  \end{aligned}$$

- Querying multiple values: If $f : D^n \to \mathcal{R}^k$, let $\Delta f = \max_{\text{neighbors } \boldsymbol{d},\boldsymbol{d'}} \|f(\boldsymbol{d}) - f(\boldsymbol{d'})\|_1$.

1.4 Count Queries Are Powerful [2]

- Special case: $D = \{0,1\}$, consider only subset-sum (count) queries $f_Q : D^n \to [n]_0$, where $Q \subseteq [n]$ and $f_Q(\boldsymbol{d}) := \sum_{i \in Q} d_i$.

- **Theorem**: Answering $O(n \log^2 n)$ randomly chosen queries with error $\mathcal{E} = o(\sqrt{n})$ allows an adversary to reconstruct most of the rows (all if $n \to \infty$) with the following algorithm:

i) [Query phase] For $j = 1, \dots, t$, choose $Q_j \subseteq [n]$ uniformly at random. Set $a_Q := K_f(\boldsymbol{d})$ where $f = f_{Q_j}$

ii) [Weeding phase] Solve linear program with unknowns $c_1, \dots, c_n$:

$$a_{Q_j} - \mathcal{E} \le \sum_{i \in Q_j} c_i \le a_{Q_j} + \mathcal{E} \qquad \forall j \in [t]$$
$$0 \le c_i \le 1 \qquad \forall i \in [n]$$

iii) [Rounding Phase] Let $c'_i = 1$ if $x_i > \frac{1}{2}$ and $c'_i = 0$ otherwise.

- Note first: LP has solution because $\boldsymbol{d}$ is feasible solution

- For $\boldsymbol{x} \in [0,1]^n$, denote by $\bar{\boldsymbol{x}}$ rounding each coordinate to the nearest multiple of $\frac{1}{n}$.

- We say $\boldsymbol{x}$ *is ε-far away from* $\boldsymbol{d}$ if $|x_i - d_i| \ge \frac{1}{3}$ for more than an ε-share of all rows i.

- If an $\boldsymbol{x}$ satisfies the LP, we have for any query Q_j selected by the algorithm

$$\Big|\sum_{i \in Q_j} (\bar{x}_i - d_i)\Big| \le \Big|\sum_{i \in Q_j} (\bar{x}_i - x_i)\Big| + \Big|\sum_{i \in Q_j} x_i - a_{Q_j}\Big| + \Big|a_{Q_j} - \sum_{i \in Q_j} d_i\Big| \le \frac{|Q_j|}{n} + \mathcal{E} + \mathcal{E} \le 1 + 2\mathcal{E}\,.$$

  Conversely, we say a query $Q \subseteq [n]$ *disqualifies* $\bar{\boldsymbol{x}}$ if $\left|\sum_{i \in Q} (\bar{x}_i - d_i)\right| > 1 + 2\mathcal{E}$.

- **Disqualifying Lemma** (without proof here, based on Azuma's inequality): Suppose $\boldsymbol{x}, \boldsymbol{d} \in [0,1]^n$, $\mathcal{E} = o(\sqrt{n})$. If $\boldsymbol{x}$ is ε-far away from $\boldsymbol{d}$ then there is $\delta > 0$ so that

$$\Pr_Q \left[\Big|\sum_{i \in Q} (x_i - d_i)\Big| > 2\mathcal{E} + 1 \right] = \delta.$$

- For any $\bar{\boldsymbol{x}}$ that is far away from $\boldsymbol{d}$, lemma says that $\bar{\boldsymbol{x}}$ is disqualified by one of the queries picked by the algorithm with probability $1 - (1-\delta)^t$. With the union bound,

$$\Pr_{Q_1,\dots,Q_t}[\forall \bar{\boldsymbol{x}} : \exists j : Q_j \text{ disqualifies } \bar{\boldsymbol{x}}] > 1 - (n+1)^n (1-\delta)^t > 1 - o(1/n^k)$$

  for any $k \in \mathbb{N}$, when choosing, say, $t = n \log^2 n$.

- Now, $\bar{\boldsymbol{c}}$ was not disqualified by any Q_j, so $\bar{\boldsymbol{c}}$ is not far away from $\boldsymbol{d}$. Hence, also $\boldsymbol{c}'$ and $\boldsymbol{d}$ differ in at most an ε-share of the rows.

1.5 Tightness

- This is tight in the following sense: If an attacker must assume that the database is random (uniform distribution), then there is a mechanism with perturbation $\mathcal{E} = \sqrt{n} \cdot (\log n)^{1+\varepsilon}$ that does reveal almost nothing by answering polynomially many queries:
- Input: Query $Q \subseteq [n]$

i) Compute $a_Q := \sum_{i \in Q} d_i$

ii) Return $\frac{|Q|}{2}$ if $|a_Q - \frac{|Q|}{2}| < \mathcal{E}$ and return a_Q otherwise

- Of course, this mechanism is useless. Remedy: Allow only sublinear number of queries [2, 5]

1.6 Non-Numerical Queries [7]

- **Definition**: Given a database $\boldsymbol{d} \in D^n$, let $q : D^n \times \mathcal{R} \to \mathbb{R}$ be a measurable scoring (weighting) function. Then define the *exponential privacy mechanism* K_f by

$$\Pr[K_f(\boldsymbol{d}) \in S] := \frac{\int_S \exp(\varepsilon q(\boldsymbol{d}, r))\, d\lambda(r)}{\int_{\mathcal{R}} \exp(\varepsilon q(\boldsymbol{d}, r))\, d\lambda(r)} . \tag{1.1}$$

(We require that q is such that the integral is bounded.)

- Define $\Delta q := \max_{r \in \mathcal{R}, \text{neighbors } \boldsymbol{d}, \boldsymbol{d}'} |q(\boldsymbol{d}, r) - q(\boldsymbol{d}', r)|$.
- **Lemma**: As defined above, K_f gives $(2\varepsilon\Delta q)$-differential privacy.
- For neighboring $\boldsymbol{d}, \boldsymbol{d}'$ the change in both numerator and denominator of (1.1) can be at most $\exp(\varepsilon\Delta q)$ each, i.e., at most $\exp(2\varepsilon\Delta q)$ in total.

1.7 Privacy as a Solution Concept for Mechanism Design [7]

- A player's strategy is said to be ε-*dominant* if no other strategy ever provides this player with more than ε additional utility.
- **Lemma**: A mechanism satisfying ε-differential privacy makes truth-telling an $(\exp(\varepsilon) - 1)$-dominant strategy for any player with a utility function that maps $\mathcal{R}$ to $[0, 1]$.
- Notation: Let $\mu_{K,f,\boldsymbol{d}}$ be the probability distribution of $K_f(\boldsymbol{d})$, i.e., $\mu_{K,f,\boldsymbol{d}}(S) = \Pr[K_f(\boldsymbol{d}) \in S]$. When unambiguous, we omit indices.
- Even stronger: Regardless of the utility function $u : \mathcal{R} \to \mathbb{R}_{\geq 0}$, no player can cause a relative change of more than $\exp(\varepsilon)$ in its utility because $\mathrm{E}[u(K_f(\boldsymbol{d}))] = \int_{\mathcal{R}} u(r)\, d\mu_{\boldsymbol{d}}(r) \leq \exp(\varepsilon) \cdot \int_{\mathcal{R}} u(r)\, d\mu_{\boldsymbol{d}'}(r) = \exp(\varepsilon) \cdot \mathrm{E}[u(K_f(\boldsymbol{d}'))]$.

1.7.1 Unlimited Supply Auctions

- Consider auctioneer with endless supply of arbitrarily divisible good. The outcome (response) is a price $p \in \mathcal{R} := [0, 1]$. Each bidder i will reveal a non-increasing demand curve $b_i : \mathcal{R} \to \mathbb{R}_{>0}$, mapping prices to desired units. Requirement: $pb_i(p) \leq 1$

- For bid vector (database) $\boldsymbol{b}$ and price p, we sell $\sum_i b_i(p)$ items, yielding revenue $q(\boldsymbol{b}, p) = p \sum_i b_i(p)$. Let OPT denote the maximum revenue.
- **Theorem**: The exponential mechanism gives 2ε-differential privacy and has expected revenue at least $OPT - \frac{3}{\varepsilon} \ln(e + \varepsilon^2 OPT m)$, where m is the number of items sold in OPT.
- Privacy follows from above lemma, as a bidder can change $q(\boldsymbol{b}, p)$ by at most $p b_i(p) \leq 1$
- Let $S_t := \{r \in \mathcal{R} \mid q(\boldsymbol{d}, r) > OPT - t\}$. Note:

$$\mu(\overline{S_{2t}}) \leq \frac{\mu(\overline{S_{2t}})}{\mu(S_t)} = \frac{\int_{\overline{S_{2t}}} \exp(\varepsilon q(\boldsymbol{d}, r))\, d\lambda(r)}{\int_{S_t} \exp(\varepsilon q(\boldsymbol{d}, r))\, d\lambda(r)} \leq \frac{\exp(\varepsilon OPT - \varepsilon 2t) \cdot \lambda(\overline{S_{2t}})}{\exp(\varepsilon OPT - \varepsilon t) \cdot \lambda(S_t)}$$

$$\left[q(\boldsymbol{d}, r) \leq OPT - 2t \text{ in numerator}, \geq OPT - t \text{ in denominator}\right]$$

$$\leq \frac{\exp(-\varepsilon t)}{\lambda(S_t)} \qquad \left[\lambda(\overline{S_{2t}}) \leq 1\right]$$

- Suppose $t \geq \frac{1}{\varepsilon} \ln\left(\frac{OPT}{t\lambda(S_t)}\right)$. Then

$$\mathrm{E}[q(\boldsymbol{d}, K_f(\boldsymbol{d}))] \geq \left(1 - \frac{\exp(-\varepsilon t)}{\lambda(S_t)}\right) \cdot (OPT - 2t) \qquad \left[\text{by previous item}\right]$$

$$= \left(1 - \frac{t}{OPT}\right) \cdot (OPT - 2t) \geq OPT - 3t$$

- Assume, w.l.o.g., that $OPT > t$ (otherwise trivial). Set $t = \frac{1}{\varepsilon} \ln(e + \varepsilon^2 OPT m)$. Since $t \geq \frac{1}{\varepsilon}$, we have $t \geq \frac{1}{\varepsilon} \ln\left(\frac{OPT m}{t^2}\right)$. Hence, by previous item, it remains to show that $\frac{t}{m} \leq \lambda(S_t)$.
- Note that for all prices $\geq OPT - \frac{t}{m}$ less than the optimal price, the same m items would continue to be sold (demand non-increasing as price increase), and for all these m items the loss is at most $\frac{t}{m}$ each. Hence, the total profit would still be at least $OPT - t$. Hence, the measure of all prices giving revenue at least $OPT - t$ (i.e., $\lambda(S_t)$) is at least as large as the measure of all prices $\geq OPT - \frac{t}{m}$ (that is, $\lambda([OPT - \frac{t}{m}, OPT]) = \frac{t}{m}$).

References

[1] T. Dalenius. Towards a methodology for statistical disclosure control. *Statistisk Tidskrift*, 15, 1977.

[2] I. Dinur and K. Nissim. Revealing information while preserving privacy. In *Proceedings of the 22nd ACM SIGMOD-SIGACT-SIGART Symposium on Principles of Database Systems (PODS'03)*, pages 202–210, 2003. DOI: `10.1145/773153.773173`.

[3] C. Dwork. Differential privacy. In *Proceedings of the 33th International Colloquium on Automata, Languages, and Programming, Part II (ICALP'06)*, volume 4052 of *LNCS*, pages 1–12, 2006. DOI: `10.1007/11787006_1`.

[4] C. Dwork. Differential privacy: A survey of results. In *Proceedings of the 5th International Conference on Theory and Applications of Models of Computation (TAMC'08)*, volume 4978 of *LNCS*, pages 1–19, 2008. DOI: `10.1007/978-3-540-79228-4_1`.

[5] C. Dwork and K. Nissim. Privacy-preserving datamining on vertically partitioned databases. In *Proceedings of the 24th Annual International Cryptology Conference (CRYPTO'04)*, pages 528–544, 2004.

[6] S. Goldwasser and S. Micali. Probabilistic encryption. *Journal of Computer and System Sciences*, 28(2): 270–299, 1984. DOI: `10.1016/0022-0000(84)90070-9`.

[7] F. McSherry and K. Talwar. Mechanism design via differential privacy. In *Proceedings of the 48th Annual Symposium on Foundations of Computer Science (FOCS'07)*, pages 94–103, 2007. DOI: `10.1109/FOCS.2007.66`.