Measures, Metrics and Meters: Some Mathematical Aspects of Privacy

Interdisciplinary Workshop on Privacy
Maynooth University Hamilton Institute

Ollie Mason
Work with Naoise Holohan & Doug Leith

September 29th, 2014

Talk Outline

1. Differential Privacy
2. Monotone Classes and Differential Privacy
3. Accuracy

Databases, Queries and Outputs

Databases

- Data entries belong to a set D (contained in some larger set U).
- Database $\mathbf{d} = (d_1, \ldots, d_n)$ in D^n.
- Each entry in , d_i corresponds to one *member* or *row*.

Queries and Outputs

- The answers/outputs of a query Q take values in some set E_Q.
- Q maps from D^n to E_Q. If $\mathbf{d}$ is the database, the correct query response is $Q(\mathbf{d})$.

Databases and Queries

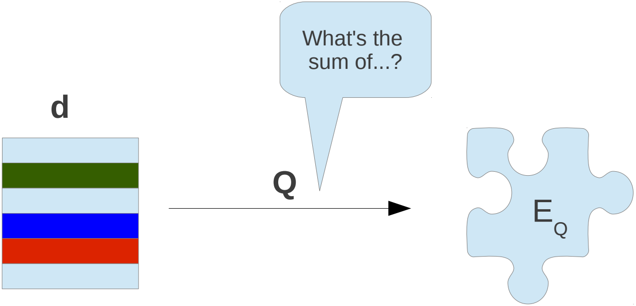

Databases, Queries and Privacy

- GENERAL AIM: Answer queries *accurately* without compromising *individual* privacy.
- A popular approach is to appropriately "*perturb*" the correct response to *a given query* on *a given database*.
- Not a new problem: Approaches developed in context of statistical disclosure control go back decades - survey article by Adam and Wortmann (1989).

Differential Privacy

- Introduced by C. Dwork in 2006 following on from earlier work by Dinur, Nissim and others (blatant non-privacy).
- CORE IDEA: If *one member* changes their entry, this does not have a “large impact” on the response to the query.

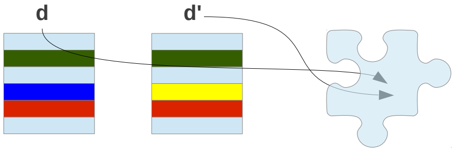

Differential Privacy

A lot of work done on different aspects:

- algorithms (eigenvalue, singular value decompositions);
- theory (lower bounds on error, optimal mechanisms, statistical implications);
- applications (Recommender Systems, Network Data).
- Focus is typically on data of some specific type and/or a particular problem.

Differential Privacy

One change doesn't make a significant difference.

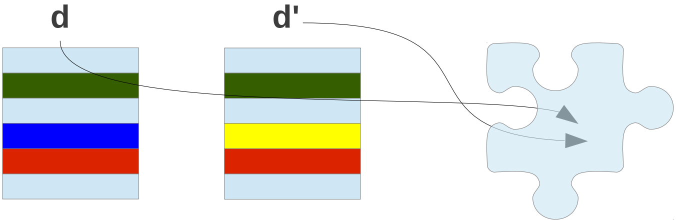

Differential Privacy in the Abstract

Advantages of taking an abstract view:

- Provides a uniform framework within which to discuss different mechanisms.
- Allows different data types to be handled in a uniform way.
- Results can be widely applied.
- Identifies precisely what is needed.
- Simplifies proofs in some instances.

Differential Privacy

- To talk about randomness or probability, we need a *Probability Space* $(\Omega, \mathcal{F}, \mathbb{P})$.
- At a minimum we need to equip the data spaces U, U^n (and naturally D, D^n), output space E_Q with *measure structures*.
- Denote by $\mathcal{A}^n$, $\mathcal{A}_Q$ the σ-algebras on U^n, E_Q.

Differential Privacy

- Given a query Q on a database , a mechanism generates a "random" response in E_Q.
- The "randomness" comes in via our probability space Ω.
- Mechanism is defined as a family of *measurable* maps $X_{Q,\mathbf{d}}$ from Ω to E_Q; we have one mapping for each $\mathbf{d}$, Q.
- $\mathbf{d} \sim \mathbf{d}'$ if they differ in exactly 1 entry.

Differential Privacy

Given $\epsilon > 0$, $\delta \geq 0$, the mechanism is (ϵ, δ) differentially private (DP) if

Differential Privacy

$$\mathbb{P}(X_{Q,\mathbf{d}} \in A) \leq e^{\epsilon}\mathbf{P}(X_{Q,\mathbf{d}'} \in A) + \delta$$

for all $\mathbf{d} \sim \mathbf{d}'$ and all $A \in \mathcal{A}_Q$.

If we set $\delta = 0$, we obtain the definition of relaxed differential privacy.

Differential Privacy

A change in one entry of **d** doesn't make a big difference to the probability of the mechanism response being in A ($p(A)$).

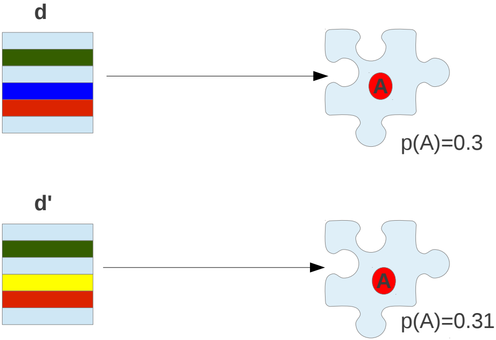

Output Perturbations

Basic Idea

Randomise the correct response to each query - usually by adding noise.

Popular approaches for real-valued data include:

- Laplacian Noise;
- Gaussian Noise

Output Perturbations

First answer the query, then perturb.

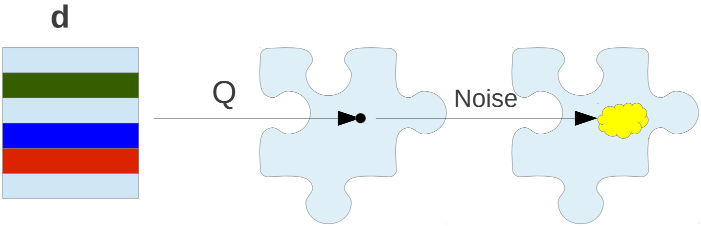

Output Perturbations - Formal Description

- For each $q \in E_Q$, $Y_q : \Omega \to E_Q$ - E_Q-valued random variable.
- If $L(\omega)$ is a Laplacian Random Variable, $Y_q(\omega) = q + L(\omega)$.
- For a query Q and database $\mathbf{d}$:

$$X_{Q,\mathbf{d}} := Y_{Q(\mathbf{d})}.$$

We do not need to assume any algebraic structure on E_Q.

Sanitised Response Mechanisms

Basic Idea

Perturb the database *once* and then answer queries on perturbed/sanitised database.

- Repeating the same query will always generate the same response.
- Intuition: Worst case is releasing the database privately.

Sanitised Response Mechanisms

First perturb the database, then answer the query.

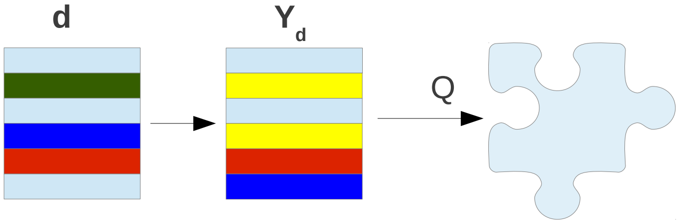

Sanitised Response Mechanisms - Formal Description

- For each $\mathbf{d}$ in D^n, $Z_{\mathbf{d}} : \Omega \to D^n$ - D^n valued Random Variable
- $Z_{\mathbf{d}}$ represents the *sanitised* database.
- For any query Q, and database $\mathbf{d}$

$$X_{Q,\mathbf{d}} := Q(Z_{\mathbf{d}}).$$

Sanitisations correspond to answering the identity query: $I : D^n \to D^n$ - $(I(\mathbf{d}) = \mathbf{d})$.

Differential Privacy and the Identity Query

The formal setup matches the intuition (phew!)

Differential Privacy for Sanitised Response Mechanisms

If the sanitisation $Z_{\mathbf{d}}$ is (ϵ, δ) differentially private then the sanitised response mechanism $Q(Z_{\mathbf{d}})$ is (ϵ, δ) differentially private for any measurable query Q.

- In the abstract setting, the argument for this is extremely simple.
- The output space E_Q could consist of sequences - in theory, we can answer an unlimited number of queries in a differentially private manner.

Sanitised Response and Output Perturbation

- Ask the same query k times: this can be modelled as a single query

$$Q^{(k)}(\mathbf{d}) = (Q(\mathbf{d}), \ldots, Q(\mathbf{d})).$$

- Result for sanitised response mechanisms guarantees that if the sanitisation is differentially private then so is the response to $Q^{(k)}$ (for any k)

Sanitised Response and Output Perturbation

In contrast, we can construct a binary-valued query and an output perturbation mechanism such that:

- the mechanism is differentially private for the query asked once;
- it violates differential privacy if the query is asked twice.

Differential Privacy and Monotone Classes

Abstract observation:

Monotonicity for Increasing Unions

If the differential privacy (DP) inequality holds for a seqence of sets $M_1, M_2, \ldots$ with $M_1 \subseteq M_2 \subseteq \ldots$ then it also holds for $M = \cup_i M_i$.

Monotonicity for Decreasing Intersections

Similarly, if the DP inequality holds for a seqence of sets $M_1, M_2, \ldots$ with $M_1 \supseteq M_2 \supseteq \ldots$ then it also holds for $M = \cap_i M_i$.

Verifying Differential Privacy

Formally, a mechanism is required to satisfy the DP inequality on all sets in the σ-algebra.

Algebra is Sufficient

If the DP inequality holds for all sets belonging to an algebra that generates $\mathcal{A}_Q$ (often a significantly smaller collection of sets), then it holds on $\mathcal{A}_Q$.

For instance, for real-valued queries it is enough to verify the inequality on finite (rather than countable) unions and intersections of intervals.

Differential Privacy and Functional Data

- Query Q takes values in $E_Q = C([0,1])$
- Can represent time-courses, continuous measurements for example.
- $C([0,1])$ equipped with Borel σ-algebra (where the topology is that given by the sup norm).
- Given a positive integer k and real numbers $0 \leq t_1 < \cdots < t_k \leq 1$ define $\pi_{t_1,\ldots,t_k} : C([0,1]) \to \mathbb{R}^k$ by

$$\pi_{t_1,\ldots,t_k}(f) = (f(t_1), \ldots, f(t_k)).$$

Differential Privacy and Functional Data

To any mechanism $X_{Q,\mathbf{d}}$ taking values in $C([0,1])$ associate the finite dimensional mechanism

$$X_{Q,\mathbf{d}}^{t_1,\dots t_k} = \pi_{t_1,\dots,t_k} \circ X_{Q,\mathbf{d}}.$$

Using the fact that DP on an algebra is sufficient, we can show:

DP for $C([0,1])$-valued queries

If the finite-dimensional mechanisms $X_{Q,\mathbf{d}}^{t_1,\dots t_k}$ are differentially private for all k and $t_1, \dots, t_k$ in $[0,1]$, then the mechanism $X_{Q,\mathbf{d}}$ is differentially private

Relates differential privacy for functional (infinite dimensional) data back to finite dimensional case.

Differential Privacy and Product Mechanisms

The monotonicity property of differential privacy also allows us to construct sanitisations coordinate-wise (entry-by-entry).

- Start with a 1-dimensional sanitisation: Y_d defined for $d \in D$.
- Define a sanitisation on D^n by $Y_{\mathbf{d}} := (Y_{\mathbf{d}}^1, \ldots, Y_{\mathbf{d}}^n)$.
- $Y_{\mathbf{d}}^i$ has the same distribution as Y_{d_i} and all the $Y_{\mathbf{d}}^i$ are independent.

Differential Privacy and Product Mechanisms

We call $Y_{\mathbf{d}}$ a *product sanitisation*.

Product Mechanisms

$Y_{\mathbf{d}}$ is differentially private if and only if Y_d (the 1-dimensional sanitisation) is differentially private.

Allows high dimensional sanitisations to be simply constructed.

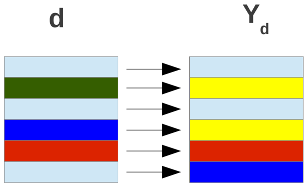

Categorical Example

- D has $m+1$ elements corresponding to categories, hobbies of individuals etc.
- Define the 1-d sanitisation

$$\mathbb{P}(Y_d = d) = 1 - pm, \quad \mathbb{P}(Y_d = d') = p,$$

  $d \neq d'$.
- Assume that $1 - pm > p$.
- Y_d is differentially private if and only if

$$p \geq \frac{1-\delta}{m + e^{\epsilon}}.$$

Accuracy for Sanitisations

- We need to measure error - assume that D is a metric space with metric (distance) ρ.
- Assume D is compact and write $\mathrm{diam}(D)$ for its diameter.
- Focus on accuracy of 1-dimensional mechanisms - can be used to obtain formulae for higher-dimensional case.
- Maximal expected error of a sanitisation Y_d

$$\mathcal{E} := \max_{d \in D} \mathbb{E}\left[\rho(Y_d, d)\right].$$

Accuracy for Sanitisations

Two results giving lower bounds for $\mathcal{E}$: one applies to general metric spaces, the other to finite metric spaces.

General Lower Bound

If Y_d is differentially private then

$$\mathcal{E} \geq (1-\delta)\left(\frac{\mathrm{diam}(D)}{2(1+e^{\epsilon})}\right).$$

Accuracy for Finite Metric Spaces

Let $\kappa = \min\{\rho(x, y) \mid x \neq y\}$.

Finite Case

Suppose D is finite containing $m + 1$ elements and Y_d is differentially private. Then:

$$\mathcal{E} \geq (1 - \delta)\left(\frac{\kappa m}{(m + e^{\epsilon})}\right).$$

Example

- Earlier example: D has $m+1$ elements; equip D with discrete metric ($\rho(d,d') = 1$ for all $d \neq d'$).
- Can define a differentially private mechanism Y_d with $p = \frac{1-\delta}{m+e^{\epsilon}}$, where $1-p = \mathbb{P}(X_d = d)$ (for all d).
- In this case $\mathcal{E}$ is given by:

$$\sum_{d' \neq d} p = mp = (1-\delta)\left(\frac{m}{m+e^{\epsilon}}\right).$$

- Lower bound is tight in this case.

Thanks!

Thank you for your attention!