Example: Suppose we have k queries of form:

q^i = # individuals in dataset D having disease i

Each query is integer-valued.

↑ counting queries

How to release the output of these k queries with ε-DP guarantee?

$$(q^1_D, q^2_D \cdots q^k_D) + (N^1 \cdots N^k)$$

↑ $N^i \sim \mathcal{L}(0, \frac{\Delta^q_1}{\varepsilon})$

What is ℓ_1-sensitivity?

$$\Delta^q_1 := \sup_{D \sim D'} \sum_{i=1}^{k} \underbrace{|q^i_D - q^i_{D'}|}_{\leq 1} \leq 1$$

But this bound is achievable because we can construct $D \& D'$ such that $q^i_D = q^i_{D'} + 1$

So: $Z_D = q_D + (N^1, \ldots, N^k)$

$N^i \sim \mathcal{L}(0, \frac{k}{\varepsilon})$

is ε-DP.

Thm, we need to add noise with per-coordinate variance of $\frac{2k}{\varepsilon^2}$ to make this query ε-DP.

Example: Let $q^1, \ldots, q^k$ be k queries of form:

$$q^i = \frac{\#\text{ individuals with some properties}}{\text{size of dataset} \leftarrow n}$$

↑ proportion queries

what's Δ_1^q ?

$$\Delta_1^q = \frac{k}{n}$$

Since this query has smaller ℓ_1-sensitivity, it's easier to be privatized than the previous query.

$\Rightarrow$ theorem says $N \sim \mathcal{L}(0, \frac{k}{n\varepsilon})$

We talked about privacy guarantee of Laplace mechanism.

But privacy makes sense only when discussing it in the privacy-utility trade-off.

To characterize this trade-off, we need to formulate utility.

Utility of Laplace mechanism

Let's begin with scalar case.

$$Z_D = q_D + N_L\left(0, \frac{\Delta^q}{\varepsilon}\right)$$

← we simply call this sensitivity (as opposed to ℓ_1-sensitivity)

we are interested in how big the gap between input & output of the mechanism is.

option 1: $E[|Z_D - q|] = E[|N|] = \frac{\Delta^q}{\varepsilon}$

this definition of utility is intuitive but doesn't offer any flexibility.

↑ prove this either directly or by invoking the fact that:

If $Z \sim L(0, b)$ then $|Z| \sim Exp(1/b)$

Option 2: $\Pr(|Z_D - q_D| \geq t)$ for some $t \geq 0$.

* we wish the output of the mechanism Z_D to be within t of the input q_D with high probability. Taking t to be 5, this means $q_D - 5 \leq Z_D \leq q_D + 5$

$$\Pr(|Z_D - q_D| \geq t) = \Pr(|N| \geq t)$$

$$= e^{-\frac{t}{\Delta t} \cdot \varepsilon}$$

Recall that we proved:

For $Z \sim L(0, b)$, we have

$$\Pr(|Z| > tb) = e^{-t}$$

This gives a precise formulation for the utility–privacy trade-off.

To have ε-DP for a scalar query with sensitivity Δ^q, we necessarily have:

$$|Z_D - q_D| \geq t$$

with probability $e^{-\frac{\varepsilon t}{\Delta^q}}$.

Given this trade-off we can answer questions like this:

what's the best privacy offered by Laplace mechanism for a query with $\Delta^q = 1$, if I tolerate

$$|Z_D - q_D| \geq 10 \quad \text{with probability } 10^{-5} \quad ?$$

Utility for Vector-Valued Laplace

Consider the query $q = (q^1, \ldots, q^k)$, & the vector-valued Laplace mechanism

$$Z_0 = q_0 + N \quad \text{where} \quad N = (N^1, \ldots, N^k)$$

& each component $N^i \sim \text{Lap}(0, \Delta_1^q / \varepsilon)$

As before, we define utility as:

$$\Pr\left(\| Z_0 - q_0 \|_1 > t \right)$$

(ℓ_1-norm)

Unfortunately, we can't compute this probability exactly. Instead we characterize an upper-bound on it.

Theorem. For any $t \geq 0$, we have:

$$\Pr\left(\| \hat{z}_D - q_D \|_1 \geq t \right) \leq K e^{-\frac{t \Sigma}{K \Delta_1 q}}$$

Proof. we can write:

$$\Pr\left(\| \hat{z}_D - q_D \|_1 \geq Kt \right) = \Pr\left(\| N \|_1 \geq Kt \right)$$

$$\leq \Pr\left(\left(\bigcap_{i=1}^{K} |N_i| < t \right)^c \right)$$

$$= \Pr\left(\bigcup_{i=1}^{K} |N_i| \geq t \right)$$

union bound $\rightarrow$

$$\leq \sum_{i=1}^{K} \Pr\left(|N_i| \geq t \right)$$

* Note that

$$\left\{ \bigcap_{i=1}^{K} |N_i| < t \right\} \subseteq \left\{ \|N\|_1 < Kt \right\}$$

$$\Rightarrow \left\{ \|N\|_1 < Kt \right\}^c \subseteq \left\{ \bigcap_{i=1}^{K} |N_i| < t \right\}^c$$

As before

$$\leq \sum_{i=1}^{K} e^{-\frac{t\varepsilon}{\Delta_i^q}}$$

So, we have:

$$\Pr\left(\| z_D - q_D \|_1 \geq Kt \right) \leq K \cdot e^{-\frac{t\Sigma}{\Delta_1^q}}$$

or equivalently:

$$\Pr\left(\| z_D - q_D \|_1 \geq t \right) \leq K e^{-\frac{t\Sigma}{K\Delta_1^q}}.$$

□

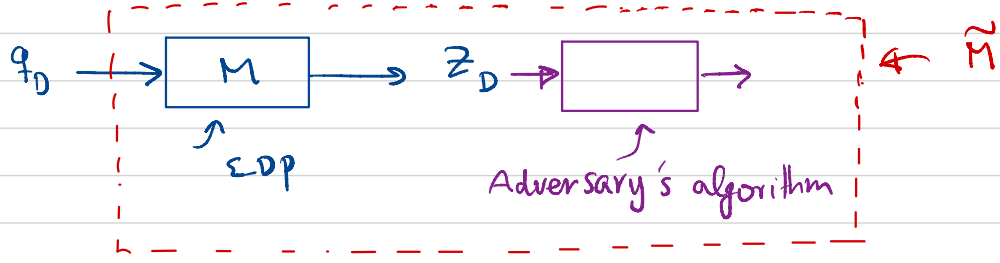
Sometimes, this utility is equivalently expressed as

$$\Pr\left(\| z_D - q_D \| \geq \frac{K\Delta_1^q}{\Sigma} \log\frac{K}{\beta} \right) \leq \beta$$

While this is not a precise privacy-utility trade-off, it has been used in practice to find achievable ε given a desired utility constraint.

Properties of ε-DP

1- Post-Processing :



Can adversary come up with an algorithm / processor to violate privacy of M? In other words, can $\tilde{M}$ be $\tilde{\varepsilon}$-DP with $\tilde{\varepsilon} > \varepsilon$?

the answer is **NO!** [Assuming adversary's algorithm only operates on the output of M & doesn't have access of dataset]

Why? Let Z_D & $Z_{D'}$ be outputs of M when running on datasets $D \sim D'$, & let Y_D & $Y_{D'}$ be the output of the adversary's algorithm.

Let $Z_D \sim M_D$ & $Z_{D'} \sim M_{D'}$

& $Y_D \sim \tilde{M}_D$ & $Z_{D'} \sim \tilde{M}_{D'}$

Since M is ε-DP, we have $E_{e^\varepsilon}(M_D \| M_{D'}) = 0 \quad \forall D \sim D'$

Then, according to DPI:

$$E_{e^\varepsilon}(\tilde{M}_D \| \tilde{M}_{D'}) \leq E_{e^\varepsilon}(M_D \| M_{D'}) = 0$$

$\Rightarrow \tilde{M}$ is ε-DP

2- Group Privacy: Does ε-DP provide privacy to a "group of individuals" as well?

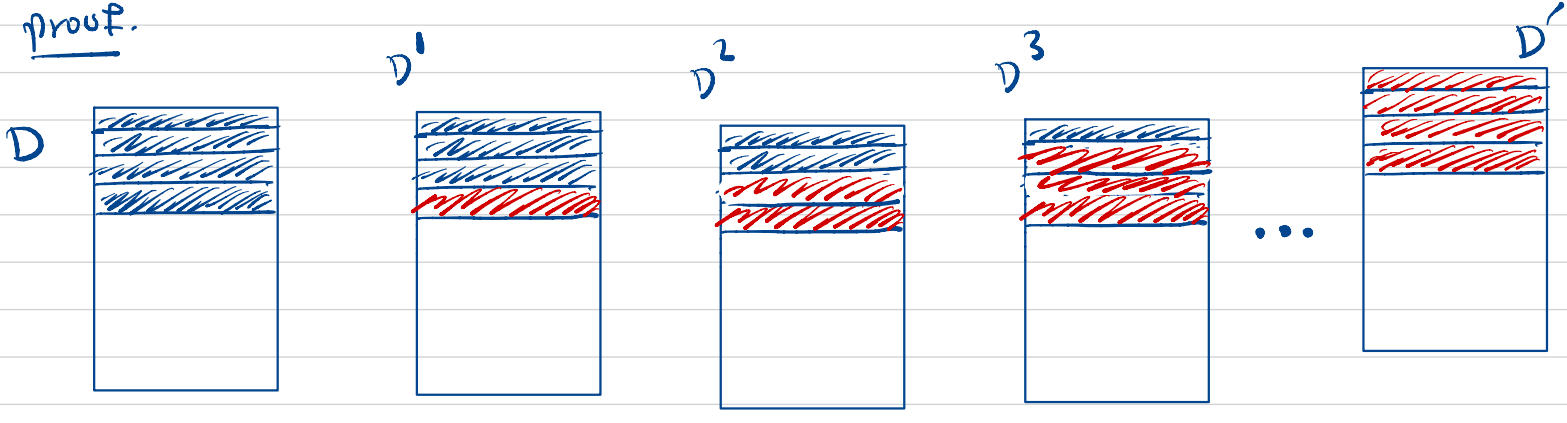
In other words, can we ensure that a "group" enjoys a plausible deniability?

theorem. Let M be an ε-DP mechanism. Then for any D & D' that differ in k entries, we have:

$$\Pr(Z_D \in E) \leq e^{k\varepsilon} \Pr(Z_{D'} \in E)$$

for any event E.

proof.



$$\Pr(Z_D \in E) \le e^{\varepsilon} \Pr(Z_{D'} \in E) \quad \text{since } D \sim D'$$

$$\le e^{2\varepsilon} \Pr(Z_{D^2} \in E) \quad \text{since } D' \sim D^2$$

$$\vdots$$

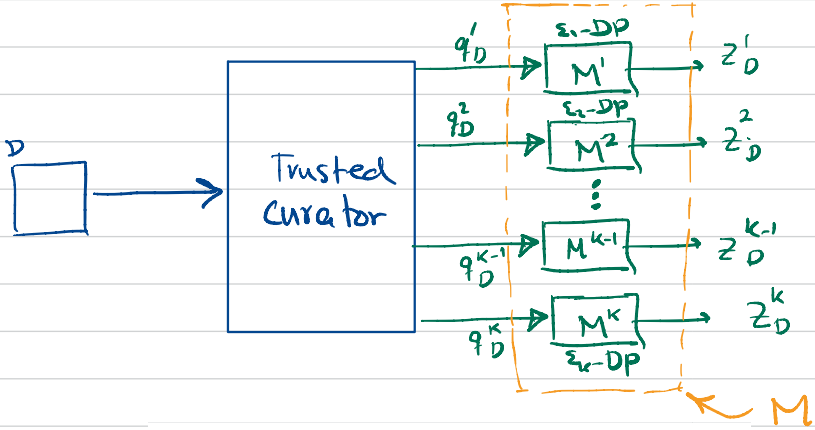
$$\le e^{(k-1)\varepsilon} \Pr(Z_{D^{k-1}} \in E) \quad \text{since } D^{k-2} \sim D^{k-1}$$

$$\le e^{k\varepsilon} \Pr(Z_{D^k} \in E) \quad \text{since } D^{k-1} \sim D'$$

∎

Composition:

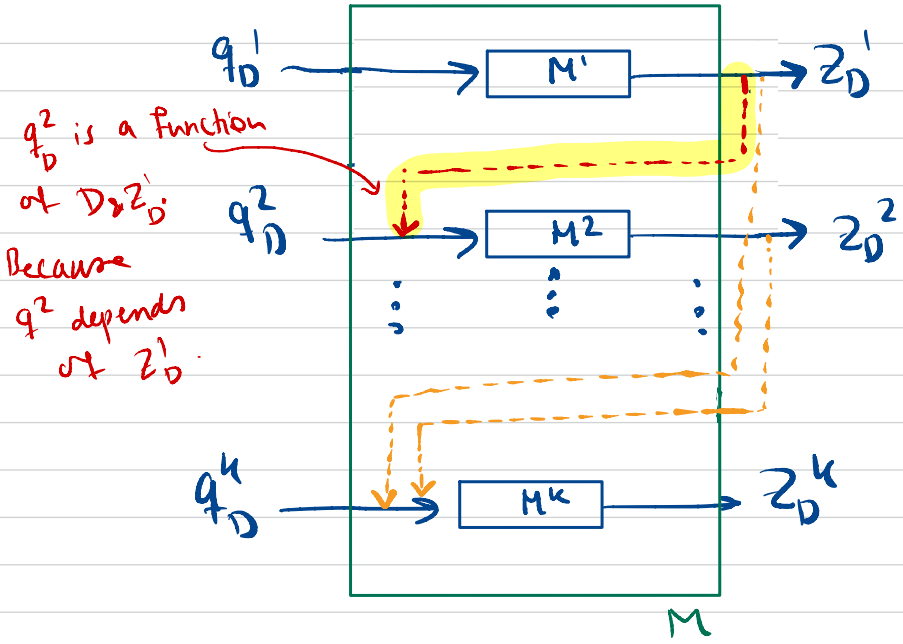
Let M^i be ε_i-DP, & we use them to answer several queries. What's the overal privacy guarantee?



what's Privacy guarantee of M?

this is an instance of non-adaptive composition, where queries are independent of each other.

In general, queries can depend on each other; a setting which's typically referred to as adaptive composition:



* If all mechanisms are the same, then their composition is usually called k-fold composition.

Theorem (Basic composition). If M^i is ε_i-DP for $i \in \{1, 2, \ldots, k\}$. Then, the composition of $M^1, \ldots, M^k$ is $\sum_{i=1}^{k} \varepsilon_i$-DP.

Proof. We give the proof only for the non-adaptive setting. The proof for the adaptive one is a bit long.

Let $Z_D^1, \ldots, Z_D^k$ be the outputs of $M^1, \ldots, M^k$, respectively.

Thus, the output of the composition mechanism is $(Z_D^1, \ldots, Z_D^k)$.

Let's find the Pdf of the output when using dataset D & D'.

$$f_{Z_D}(\vec{x}) = \prod_{i=1}^{k} f_{Z_D^i}(x_i)$$

since all z^i are independent

$$f_{Z_{D'}}(\vec{x}) = \prod_{i=1}^{k} f_{Z_{D'}^i}(x_i)$$

since each M^i is ε_i-DP

thus,

$$\frac{f_{Z_D}(\vec{x})}{f_{Z_{D'}}(\vec{x})} = \prod_{i=1}^{k} \frac{f_{Z_D^i}(x_i)}{f_{Z_{D'}^i}(x_i)} \leq \prod_{i=1}^{k} e^{\varepsilon_i} = e^{\sum \varepsilon_i}$$

∎

Composition results, such as basic composition, are essential for designing private ML. In such cases, an important concept is privacy concept.

Answering k non-adaptive queries

Vector-valued Laplace:

$$\vec{Z}_D = (q_D^1, \ldots, q_D^k) + (N^1, \ldots, N^k)$$

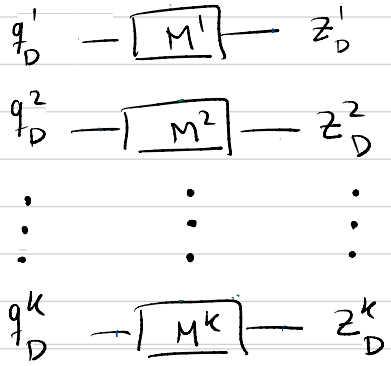
$N^i \overset{iid}{\sim} \mathcal{L}(0, b)$

* This mechanism is ε-DP if $b = \frac{\Delta_1^q}{\varepsilon}$.

Suppose each q^i is counting.

$\Rightarrow \Delta_1^q = k$

Thus, per-component noise is $\mathcal{L}(0, \frac{k}{\varepsilon})$.



Privacy budget $= \varepsilon$

If each mechanism is ε'-DP, then

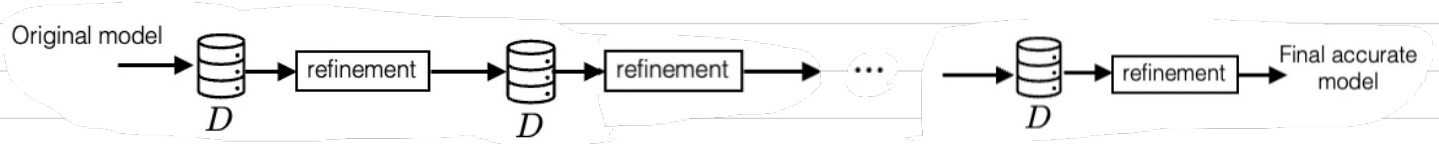
basic compo $\Rightarrow$ $\underbrace{k\varepsilon'}_{=\varepsilon}$-DP

we want each mechanism to be $\frac{\varepsilon}{k}$-DP

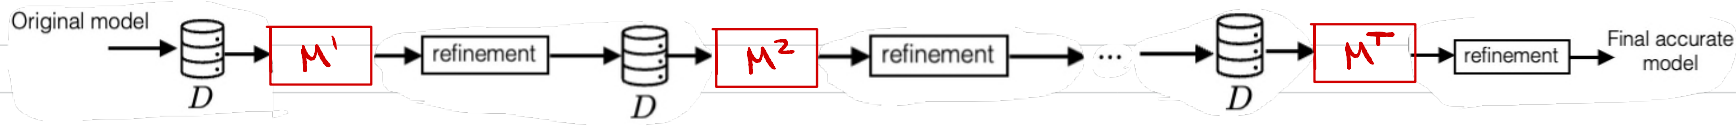
Each mechanism adds $\mathcal{L}(0, k/\varepsilon)$

Privacy budget

Any ML training algorithm can be viewed as an iterative process, in each of iteration dataset is accessed once.



To make this algorithm differentially private, we need to pass any computation on dataset through a mechanism:



Thus, all training algorithm can be thought of as adaptive composition.

If each mechanism M^i is Σ-DP, then according to basic composition, this iterative algorithm is $T\Sigma$-DP, after T number of iterations.

Usually, in practice, we know the desired level of privacy; which is typically termed privacy budget.

Given the privacy budget Σ^* for the above algorithm, we have from basic composition: $\sum_{i=1}^{k} \Sigma_i = \Sigma^*$

If all mechanisms are the same (say Σ-DP), then we must have

$$\varepsilon = \frac{\varepsilon^*}{T}$$

Now, we need to design mechanism for each iteration that is $\frac{\varepsilon^*}{T}$-DP.

Remark: If we want to make an algorithm to be, say, 5-DP, then each iteration should be $\frac{5}{T}$-DP. The issue is that T is often very large ($\approx 10^6$). Thus, each iteration must be extremely private, or equivalently, the noise is dominant $\Rightarrow$ the algorithm can't be accurate!

This is caused because of basic composition, & would be improved if we can come up with a better result than basic composition.

Question: Is basic composition optimal? that is, is there an ε-DP mechanism M such that its k-fold composition is ε^*-DP, where $\varepsilon^* < k\varepsilon$?

Basic composition is **optimal**!

Example: Consider k queries $q^1, \ldots, q^k$, each of which is counting.

"so sentivity = 1"

Suppose D & D' are neighboring & $q_D^j = q_{D'}^j + 1 \quad \forall j$.

Let Z_D & $Z_{D'}$ be the k-dimensional output of a k-fold composition of a Laplace mechanism & let f_{Z_D} & $f_{Z_{D'}}$ denote their pdf. For any $x \in \mathbb{R}^k$

$$\frac{f_{Z_D}(x)}{f_{Z_{D'}}(x)} = \prod_{i=1}^{k} e^{\varepsilon |x_i - q_{D'}^i| - \varepsilon |x_i - q_D^i|}$$

Let $x = (a, a, \ldots, a) \in \mathbb{R}^k$ where $a > q_D^j \quad \forall j$.

thus:

$$\frac{f_{Z_D}(x)}{f_{Z'_D}(x)} = \prod e^{\varepsilon(q^i_D - q^{i'}_D)} = e^{k\varepsilon}$$

Thus, the ratio of pdf is exactly equal to $e^{k\varepsilon}$, & hence there exist at least one event E such that

$$\frac{\Pr(Z_D \in E)}{\Pr(Z'_D \in E)} = e^{k\varepsilon},$$

implying basic composition can't be improved in general.

[Note that this doesn't mean that basic composition can't be improved for a particular mechanism]

this begs the question:

How can we integrate DP into accurate ML?

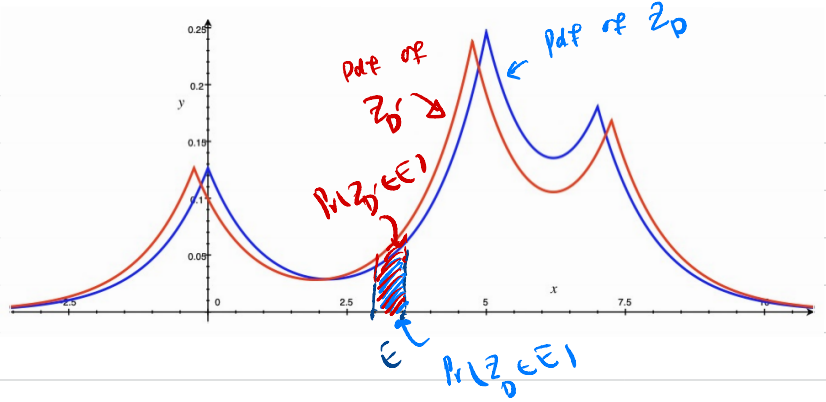
Pure DP turns out to be very stringent, meaning it's too strong to be applicable in practice!

Recall that:

A mechanism M is ε-DP if

$$e^{-\varepsilon} \leq \frac{Pr(Z_D \in E)}{Pr(Z_{D'} \in E)} \leq e^{\varepsilon}$$ for any possible event E

$$e^{-\varepsilon} \leq \frac{\text{Area under Blue}}{\text{Area under red}} \leq e^{\varepsilon}$$



Two pdf's should be within e^{ε} of each other.

What it means is that:

if there exists an event E impossible to occur for D'

then, it has to be impossible for D too.

If $Pr(Z_{D'} \in E) = 0 \Rightarrow$ Then the ratio $\frac{Pr(Z_D \in E)}{Pr(Z_{D'} \in E)}$ becomes ∞ if $Pr(Z_D \in E) > 0$

$\Rightarrow \varepsilon = \infty$ i.e., NO Privacy guarantee

So: $Pr(Z_{D'} \in E) = 0 \iff Pr(Z_D \in E) = 0$

This is very stringent. Why?

Suppose there exist E: $\Pr(Z_D \in E) = 0$

$\Pr(Z_{D'} \in E) = 10^{-7}$

* So NO privacy guarantee even though E is extremely unlikely!

E is a "bad" event: because it reveals something about dataset!

The definition of ε-DP is unpractically stringent