Friday Live Exercises

Privacy-preserving Data Release 2

How to achieve Differential Privacy Output Perturbation

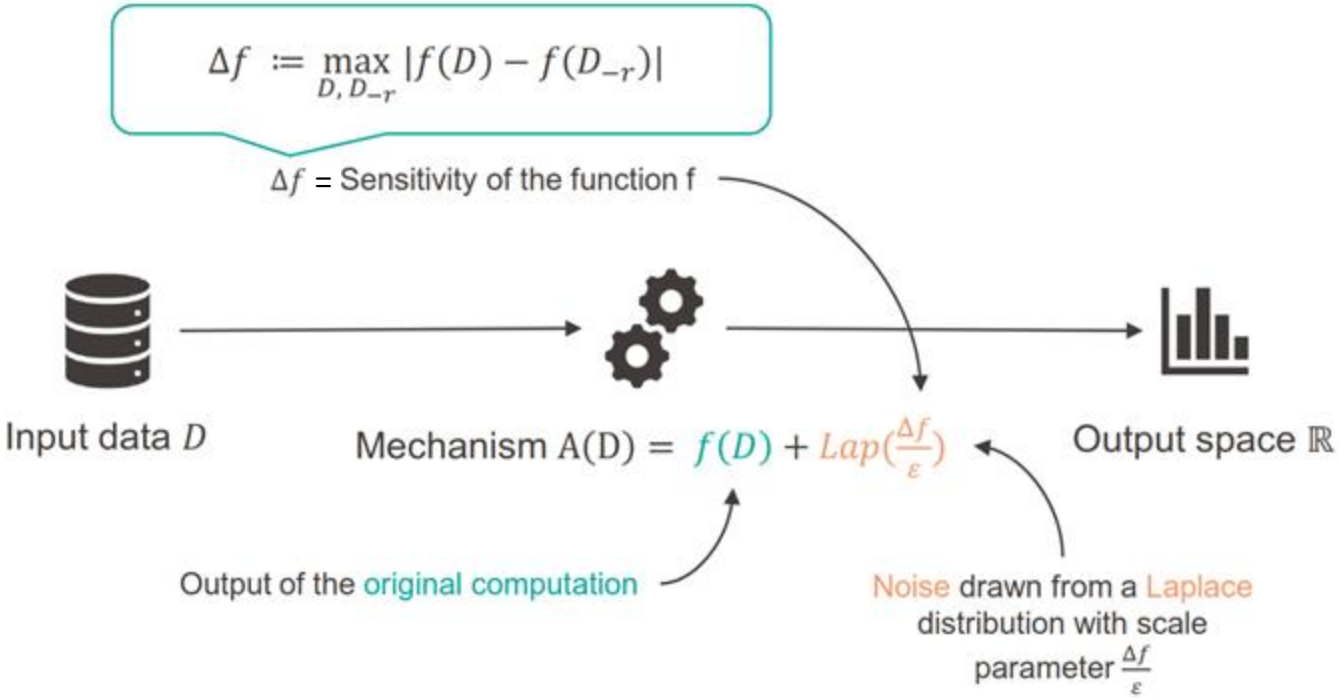

Theoretical Warmup: Sensitivity 1

	name	secret
1	Alice	29
2	Bob	2
3	Charlie	35
4	Dylan	29
5	Eric	34
6	Frank	12

The sensitivity of the query

`MAX(secret) SELECT(1,2,4,6)`

for value `secret` that takes values in `[2, 40]` over the database shown on the left is:

a) 27

b) 38

c) 33

d) 40

Theoretical Warmup: Sensitivity 2

	name	day	secret
1	Alice	Mon	29
2	Bob	Tue	2
3	Charlie	Wed	35
4	Alice	Tue	10
5	Eric	Mon	34
6	Bob	Wed	12

The sensitivity of the query

```
SUM(secret) GROUPBY day
```

for value `secret` that takes values in `[2, 40]` over the database shown on the left is:

a) 40

b) 120

c) 38

d) 280

Waterwolf

The Waterwolf Browser company collects usage statistics from its users to better understand which websites its users visit most frequently and how this behaviour changes over time. In the Waterwolf database, each user is identified by a unique identifier and the **`usage_table`** contains the following information about each user:

- the origin country of their IP address
- their total browsing time in minutes capped at a maximum of 1000
- a list of binary values that indicates for a pre-defined set of 1000 websites whether the user has ever visited this website

The Waterwolf database gets updated with the most up-to-date statistics on a daily basis. This means that when a new user has started using the Waterwolf Browser, a new entry is created; and that when an existing user visits a website they had not previously visited, the corresponding entry is flipped from 0 to 1.

user_id	country	usage_time	google.com	amazon.com	…	protonmail.com
uid198	CH	121	0	0	…	1
uid847	CH	76	1	1	…	0
…	…	…	…	…	…	…
uid272	FR	876	1	0	…	1

Waterwolf

Daria is a data analyst at Waterwolf tasked with analysing user behaviour. Daria does not have direct access to the database, but she can run analysis scripts on its content. She has written an analysis script that gives her for each of the 1000 pre-defined websites the number of users per country that have visited this website and the total browsing time of all users in a given country. The pseudocode of her analysis script is shown below. Daria runs the script once on the current usage database at the start of every day:

```
FOR website in website_list:
     SELECT country, SUM(website) FROM usage_table GROUPBY country

SELECT country, SUM(usage_time) FROM usage_table GROUPBY country
```

Part 1: Daria has finally convinced her friend Alfredo to start using the Waterwolf Browser. This is a great success for Waterwolf as Alfredo is the first Waterwolf user in Italy. He starts browsing the web on Monday and at the end of the day sends his friend Daria an excited message: “Waterwolf is really great! But I am a bit concerned that you can now learn everything about my browsing habits.”

Is Alfredo right to be concerned? Justify your answer and, if you think there is reason for concern, explain what information Daria might learn about Alfredo and how she might learn it. Recall that Daria does not have direct access to the usage database. She can only observe the output of her analysis script at the start of each day.

Waterwolf

Part 2: Waterwolf gets contacted by privacy researchers who are concerned about the privacy risks of collecting usage statistics. The researchers recommend Daria to use the differential privacy model to reduce the privacy leakage of her analysis script. According to the researchers, Daria's analysis should not exceed a total privacy budget of epsilon=2 for each user over the course of a week.

Describe a differentially private version of Daria's analysis script. Your description needs to include details on what noise addition mechanism Daria should use, how she needs to scale the noise, and how the noise gets added to the results. Argue why your proposed algorithm achieves user-level differential privacy with a total epsilon of 2 after a week of analysis. Hint: Apply the sequential and parallel composition theorems.