Friday Live Exercises

Privacy-preserving Data Release 1

Race	ZIP
Asian	02138
Asian	02139
Asian	02141
Asian	02142
Black	02138
Black	02139
Black	02141
Black	02142
White	02138
White	02139
White	02141
White	02142

Original

Race	ZIP
*	02138
*	02139
*	02141
*	02142
*	02138
*	02139
*	02141
*	02142
*	02138
*	02139
*	02141
*	02142

Release 1

Race	ZIP
Asian	0213*
Asian	0213*
Asian	0214*
Asian	0214*
Black	0213*
Black	0213*
Black	0214*
Black	0214*
White	0213*
White	0213*
White	0214*
White	0214*

Release 2

Part 1: PrivateData, Inc. wants to release a sorted dataset of voter registrations to Client *(Only the quasi-identifier columns are shown).*

1. The Client is interested in the ZIP code distribution in the dataset. To enable such analysis, PrivateData sanitizes the Race column to get k-anonymity and produces the table in Release 1.
2. The week after, the Client changes their mind and decides that they are interested in voter demographics. PrivateData releases the second release of the dataset which hides the last digit of the ZIP code.

What are the k-anonymity parameter of the two released tables considered individually? Can the Client reconstruct the original dataset after the second release? If yes, how to mitigate this?

This one's easy.
k-anonymity: k = 3 **for release 1**,
k = 2 for release 2.
Attack: If the datasets are released in any order, including simultaneously, the client can join and obtain the original dataset.
Defence: If the rows of the released datasets are shuffled, then this attack is not possible.

Takeaway: Do not forget to shuffle rows in complementary/sequential releases of the same dataset.

	Race	Birth Date	Gender	ZIP	Symptom
1	black	1965	M	02141	short of breath
2	black	1965	M	02141	chest pain
3	*	1965	F	0213*	painful eye
4	*	1965	F	0213*	wheezing
5	black	1964	F	02138	obesity
6	black	1964	F	02138	chest pain
7	white	1964	M	0213*	short of breath
8	*	1965	F	0213*	hypertension
9	white	1964	M	0213*	obesity
10	white	1964	M	0213*	fever
11	white	1967	M	02138	vomiting
12	white	1967	M	02138	back pain

Release Week 1

Race	Birth Date	Gender	ZIP	Symptom
black	1965	M	02141	short of breath
black	1965	M	02141	chest pain
black	1965	F	02138	painful eye
black	1965	F	02138	wheezing
black	1964	F	02138	obesity
black	1964	F	02138	chest pain
white	196*	M	02138	short of breath
white	196*	*	02139	hypertension
white	196*	*	02139	obesity
white	196*	*	02139	fever
white	196*	M	02138	vomiting
white	196*	M	02138	back pain
black	1965	M	02139	headache
black	1965	M	02139	rash

Release Week 2

Part 2: PrivateData also sells medical data to the Client. The Client and PrivateData agree that PrivateData will release a k-anonymized version of the medical data every week. The two tables above show the data PrivateData releases in the first and second week.
Assuming that PrivateData randomly shuffles the order of rows before each release:

- What is the k-anonymity parameter with respect to *QI = {Race, BirthDate, Gender, ZIP}* of Release 1?
- What is the k-anonymity parameter with respect to *QI = {Race, BirthDate, Gender, ZIP}* of Release 2?
- Does publishing Release 2 introduce any new privacy risks for the records in the original dataset? If yes, how could you mitigate these new risks?

Each release taken separately is k-anonymous with respect to QIs={Race, BirthDate, Gender, ZIP} with k=2

Yes, the second release increases the privacy risk of records in the original dataset.

Because Symptom has not been considered part of the QI set, an adversary without any further background knowledge can reconstruct the original records through linking the Release 1 and Release 2 tables. Linkage here is the adversarial strategy (or attack). The resulting privacy loss is that we have reconstructed the original dataset: we have restored the original values in rows containing sanitized values (3–4, 7–12). This is not the same as *re-identification* which could be another type of privacy threat. As a next step, an adversary that has the reconstructed data and some background knowledge on some of the records (e.g., has an actual person's identity linked to QI = {Race, BirthDate, Gender, ZIP}) could launch a re-identification attack to link an identity to a record. This could then further lead to inference of a sensitive attribute (Symptom).

Defence: Instead of creating a new sanitization scheme for Release 2, PrivateData should add the new k-anonymized rows (red) to the Release 1 table.

Takeaway: You should **not switch** your sanitisation strategy throughout. If you release two different sanitised versions of the same data you need to consider the intersection between the version when you assess the risk.

Become A Chef

BecomeAChef is a website that publishes recipes. They distinguish themselves from other sites by having recipes tailored to people with dietary restrictions. To win new users, their marketing team launches a campaign in popular newspapers. They decide to publish in each newspaper the following pseudonymised table about the recipes viewed by users during one week:

| userID | recipe | dietary_restrictions | day| visit_duration | comments_on_recipe |

userID: random ID assigned by the system to a user viewing the page, persistent across views
recipe: name of the recipe
dietary_restrictions: typical dietary restrictions that can eat this recipe safely (e.g., celiac, halal, vegan)
day: day of the week in which userID visited the recipe's webpage
visit_duration: amount of time userID spent on the recipe's webpage
comments_on_recipe: text of the comments the user left on the recipe's webpage (blank if no comment)

The marketing team hopes that curious readers will visit their website to read more about their recipes.

Part 1: Describe a privacy concern related to the publication of this table, and an adversary (background information, capabilities, and goal) under which that privacy concern would materialise.

Part 2: Propose a defence based on k-anonymity that would mitigate the risk identified in Part 1. Justify your choice of quasi-identifiers, and the mechanisms you would use to achieve k-anonymity on each of them.

Part 1 example solution:
Concern: Alice has a particular dietary restriction that is revealed after publication.
Adversary: Alice's workmate that sees her visiting the BecomeAChef several days at the office.

Part 2 example solution:
Generalization on visit duration/UserID and suppression on day of the week.
UserID, day and duration is what the colleague uses in part 1 to single out Alice.
The goal is to increase the size of the anonymity set. (e.g. every user that visits a recipe the same number of times as Alice).