Differentially Private Data Analysis

Jonathan Ullman

Integrating Approaches to Privacy Across the Research Lifecycle
September 24-25, 2013
Harvard University

Outline

- What is differential privacy?
- What can we do with differential privacy?
- What should we talk about?

Differential Privacy (DP) [DN03,DN04,BDMN05,DMNS06,D06]

- Definitional approach to privacy
  - *Ad-hoc* methods often leads to breaches
  - Definitions can help delineate what is and is not possible
- Differential privacy is a strong, robust, flexible privacy definition
  - Intuition: "My data should not influence the outcome of the study (much)."

Differential Privacy (DP) [DN03,DN04,BDMN05,DMNS06,D06]

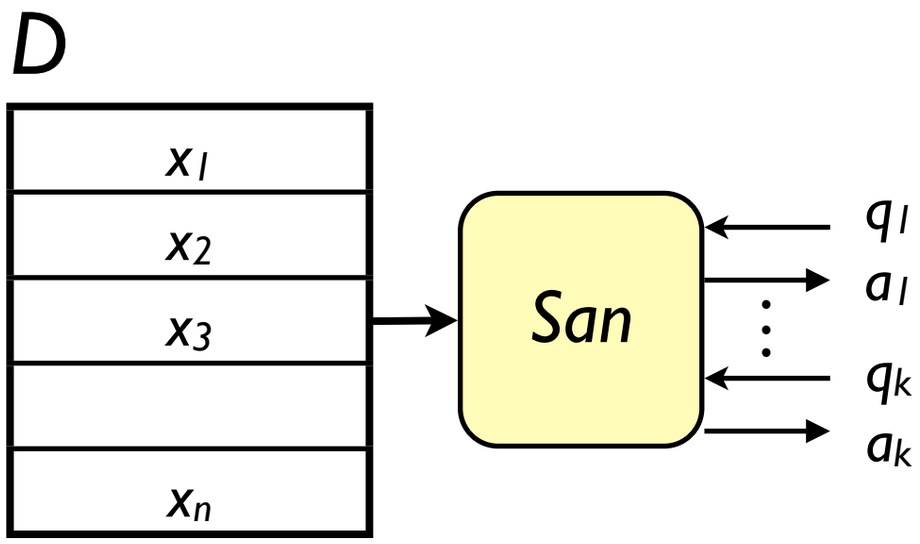

Differential Privacy (DP) [DN03,DN04,BDMN05,DMNS06,D06]

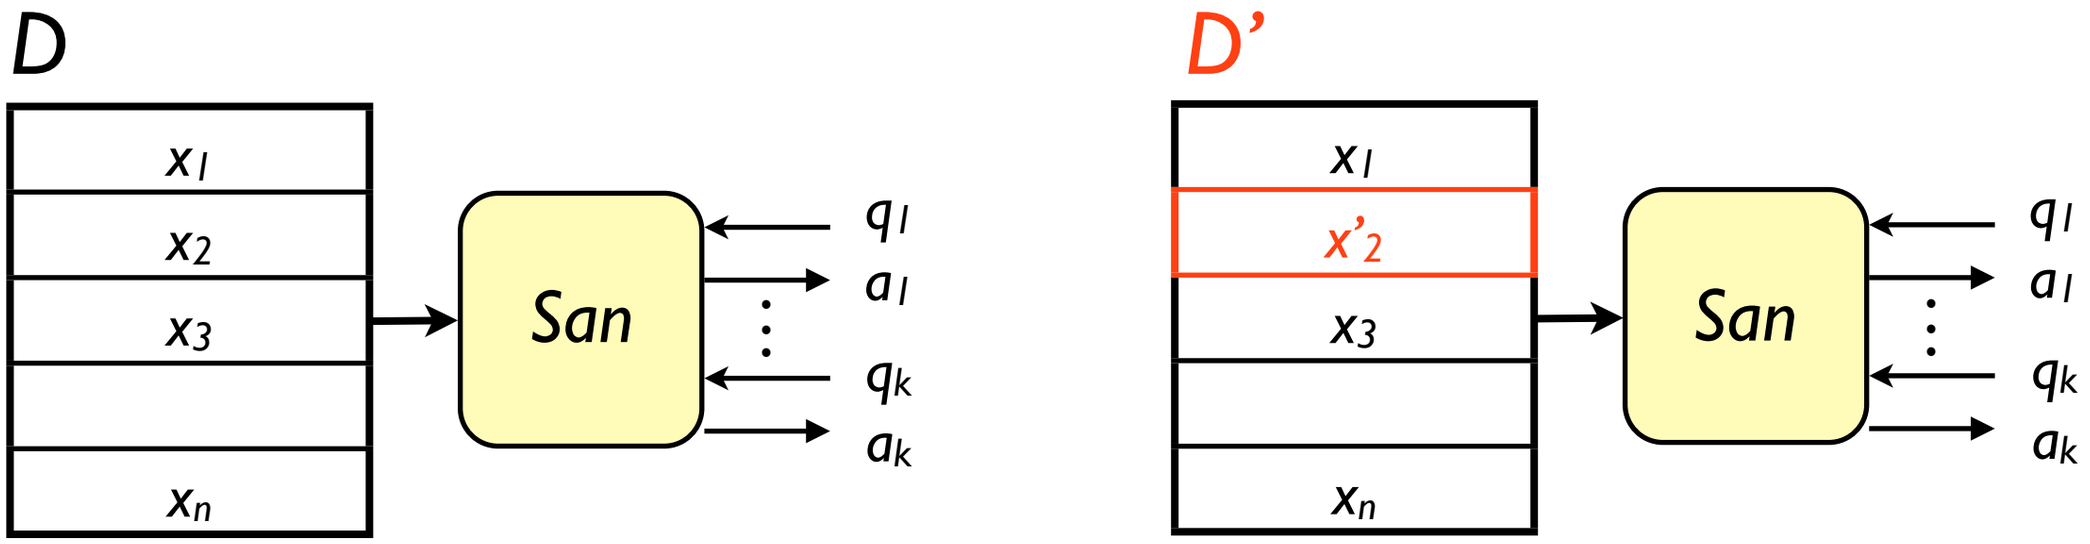

D and *D'* are neighbors if they differ only on one user's data

Differential Privacy (DP) [DN03,DN04,BDMN05,DMNS06,D06]

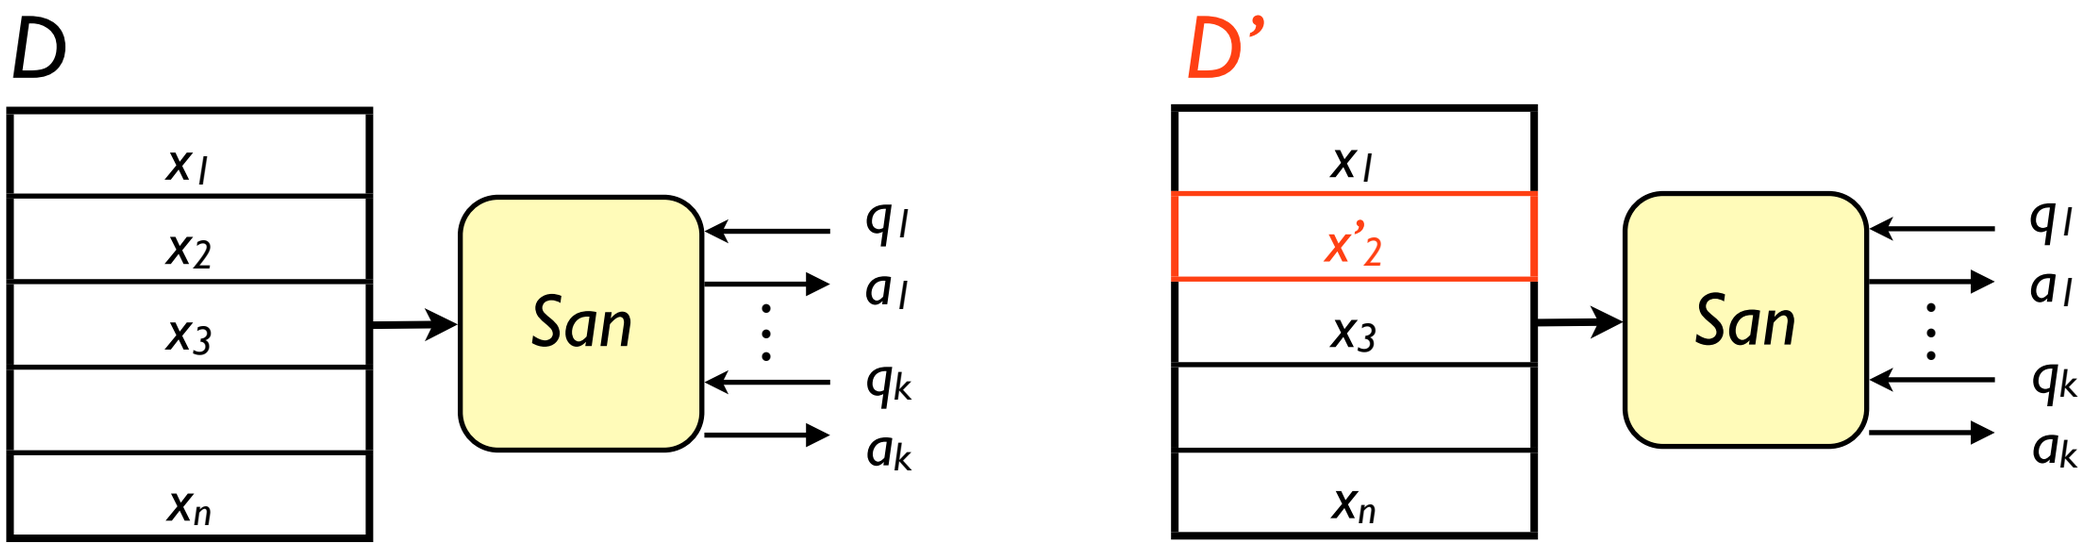

D and *D'* are neighbors if they differ only on one user's data

Definition: A randomized *San* is differentially private if for all neighbors *D*, *D'*

$$San(D;q_1,\ldots,q_k) \approx San(D';q_1,\ldots,q_k)$$

Close as distributions

Differential Privacy (DP) [DN03,DN04,BDMN05,DMNS06,D06]

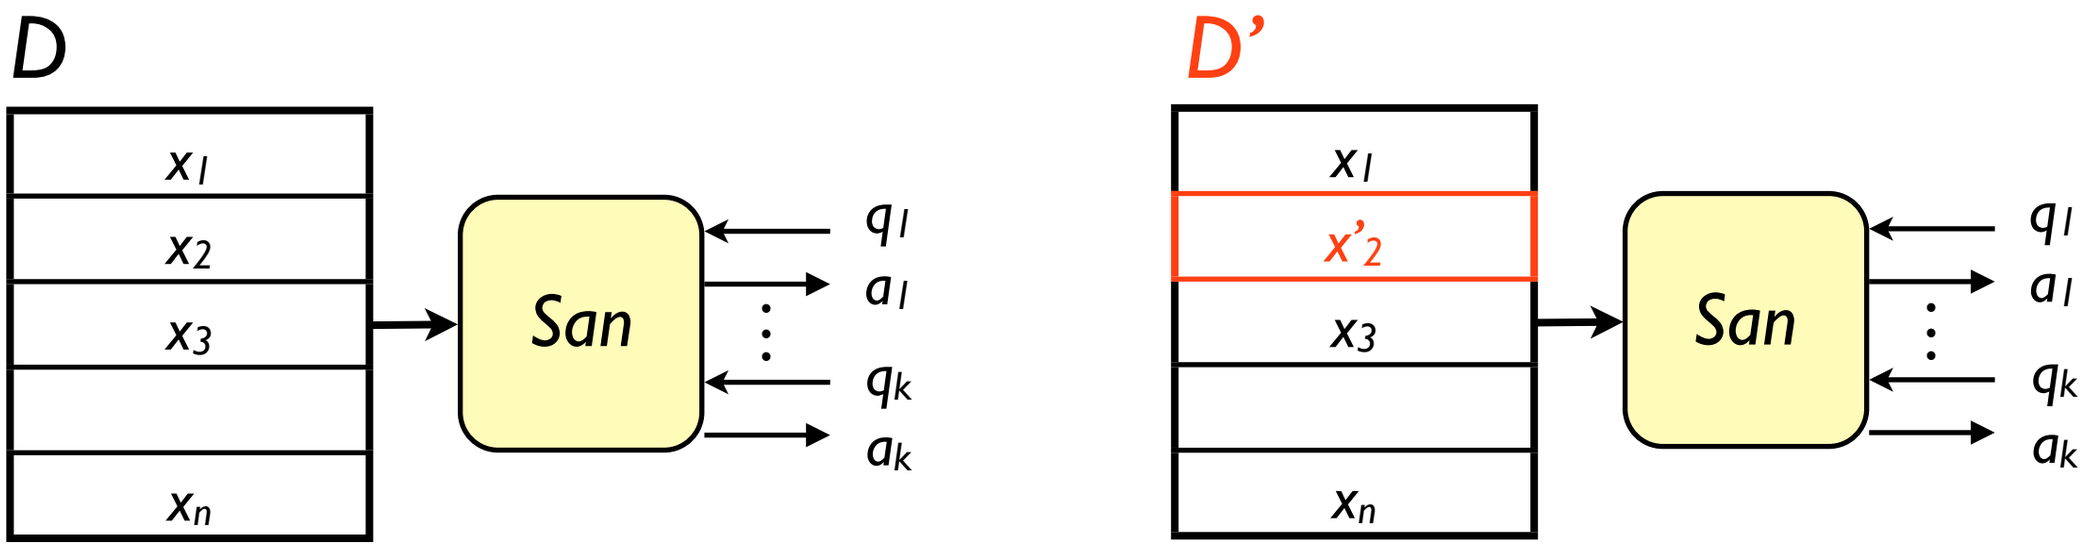

D and *D'* are neighbors if they differ only on one user's data

Definition: A randomized *San* is ε-differentially private if for all neighbors *D*, *D'*

$$San(D;q_1,\ldots,q_k) \approx_\varepsilon San(D';q_1,\ldots,q_k)$$

Close as distributions

Differential Privacy (DP) [DN03,DN04,BDMN05,DMNS06,D06]

- A DP algorithm reveals "little" information about your record in the database
- Does not mean: "I can't learn much about you from seeing the output"
  - Example: I know you're two inches taller than average, I learn the average height, so I learn your height
  - Unavoidable in general [DN10]
- Does mean: "I can't learn much about you from seeing the output because you were in *D*"

Differential Privacy (DP) [DN03,DN04,BDMN05,DMNS06,D06]

- Differential privacy quantifies loss of privacy, algorithms are not simply private or non-private
  - Makes it possible to analyze composition of differential privacy: running *10* algorithms with ε-DP satisfies *10*ε-DP
  - Makes it possible to build complicated algorithms out of simpler algorithms

What can we do with DP?

- For *San(D,q)* and *San(D',q)* to be close distributions, *San* must add noise to hide the contribution of any individual. Answers cannot be exact!
  - In statistical analyses on large databases, individuals should have only a small contribution, so noise will be small
  - If individuals can have a large effect on the analysis (e.g. outliers) then it will be necessary to add a lot of noise
  - In particular, you can't see the raw data
  - If noise is too small, then the database can be reconstructed from the answers [DN03,...]

What can we do with DP?

- For *San(D,q)* and *San(D',q)* to be close distributions, *San* must add noise to hide the contribution of any individual. Answers cannot be exact!
  - In statistical analyses on large databases, individuals should have only a small contribution, so noise will be small

$$San(D;q) = q(D) + Noise$$

- $|Noise| \to 0$ *as* $n \to \infty$
- How big *n* has to be for noise to be manageable depends on the context and the state-of-the-art algorithms

What can we do with DP?

- Even though the definition is strong, many powerful and elegant algorithms are known for DP data analysis.

- histograms [DMNS06]
- contingency tables [BCDKMT07,GHRU11,CKKL12,HRS12,TUV12,CPSY13,CTUW13]
- PAC learning [BDMN05,KLNRS07]
- machine learning (e.g. regression) [DL09,CMS11,S11,KST11,ST12,ST13,JT13]
- clustering [BDMN05,NRS07]
- social network analysis [HLMJ09,KRSY11,GRU12,KNRS13,BBDS13]
- approximation algorithms [GLMRT10]
- linear programming [HRRU13]
- singular value decomposition / low-rank approximation [HR12,KMT13,HR13]
- streaming algorithms [DNRY10,DNPR10,MMNW13]
- .....

Questions for the Workshop

- Data that isn't "row structured"?

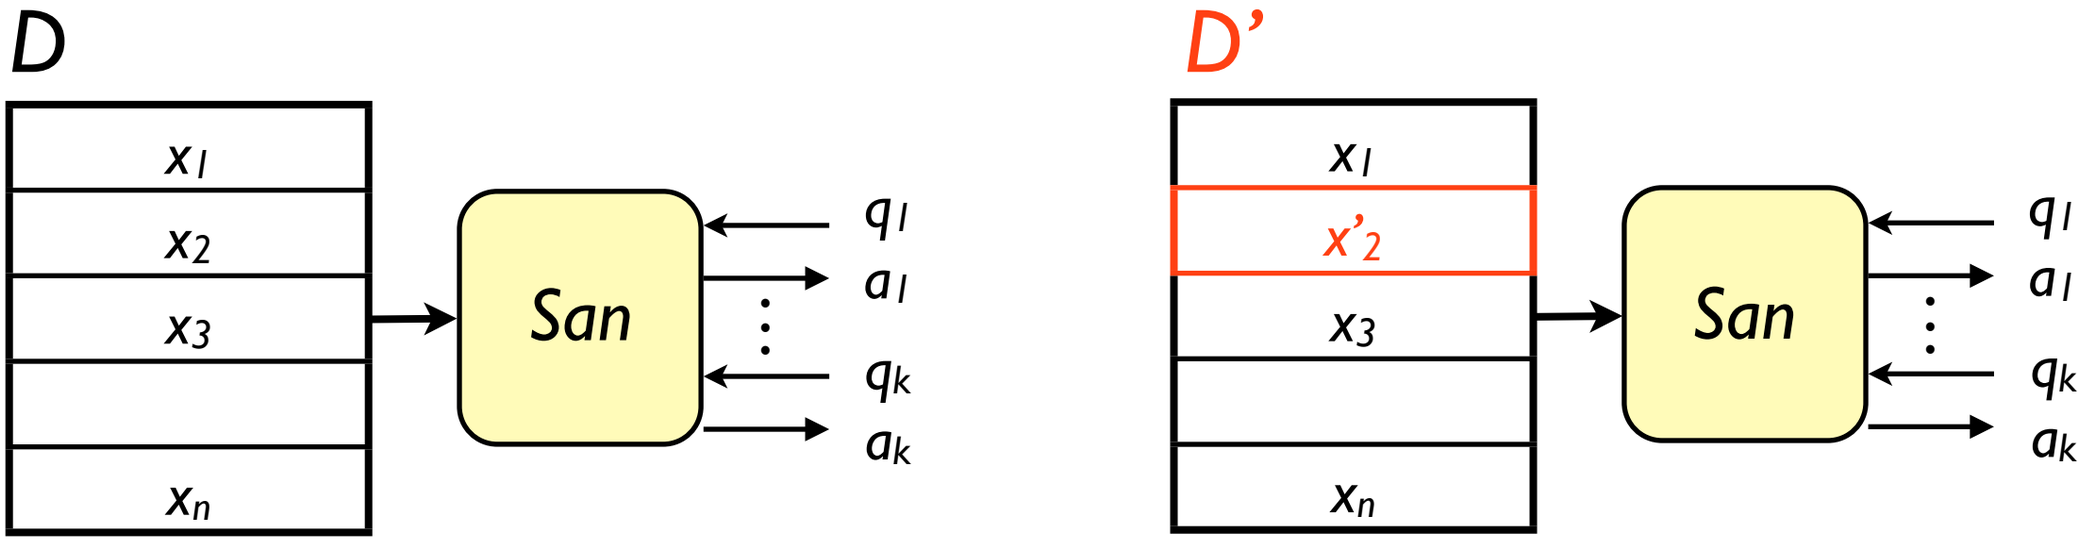

- Requiring that *San(D;q)*, *San(D';q)* be close distributions was supposed to capture "hiding information about one individual"
  - Only true if information about an individual is confined to one (or a few) rows

Questions for the Workshop

- How does differentially private data analysis fit with the current practice of data analysis?
  - As opposed to "anonymized datasets," most DP algorithms mediate interaction between the analyst and the database, and don't allow access to the raw data

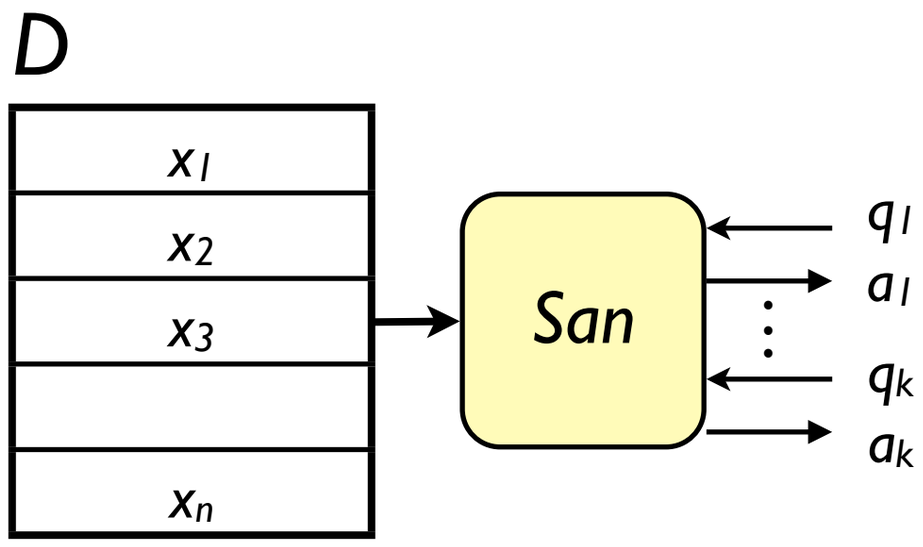

- When is touching the data essential?

RIVATE ECONOMY LABOR QUALITY, AND UNDERLYING MATRICES

< Back to Study

:1902.1/OYSLSQBRJPUNF:3:lWnju7EDKlloCWqKzdb3ig==

rsion: 1– Released: Wed Nov 28 00:00:00 EST 2007

ta File: datafile.tab

Download Subset | Recode & Case-Subset | Descriptive Statistics | **ADVANCED STATISTICAL ANALYSIS**

Selected Variables

Logistic Reg for Binary Dep Vars

More Information about the Model

Dependent

> sex

<

Output Options

- [x] Include Summary Statistics
- [x] Include Plot
- [x] Include Replication Data

Explanatory

>

<

class
age
ed2hour
ed1hour

Analysis Options

- [ ] **Simulations**

Run Model

Dataverse Analysis

The following are the results of your requested analysis.

Summary Results

- Call: zelig(formula = sex ~ class + age + ed1hour + ed2hour, model = "logit", data = data)

Deviance Residuals:

Min	1Q	Median	3Q	Max
-8.4904	0.0000	0.0000	0.0001	8.4904

Coefficients:

	Estimate	Std. Error	z value	Pr(>|z|)
(Intercept)	2.0761e+13	2.5442e+13	0.8160	0.4145
class	5.9152e-03	3.9310e-01	0.0150	0.9880
age	-2.0761e+13	2.5442e+13	-0.8160	0.4145
ed1hour10012835	4.1522e+13	5.0883e+13	0.8160	0.4145
ed1hour100285552	8.3044e+13	1.0177e+14	0.8160	0.4145
ed1hour1004600704	6.2283e+13	7.6325e+13	0.8160	0.4145
ed1hour100926200	6.2283e+13	7.6325e+13	0.8160	0.4145
ed1hour1011177792	1.0381e+14	1.2721e+14	0.8160	0.4145
ed1hour1011535104	1.0381e+14	1.2721e+14	0.8160	0.4145

Questions for the Workshop

- How does differentially private data analysis fit with the current practice of data analysis?
  - As opposed to "anonymized datasets," most DP algorithms mediate interaction between the analyst and the database, and don't allow access to the raw data

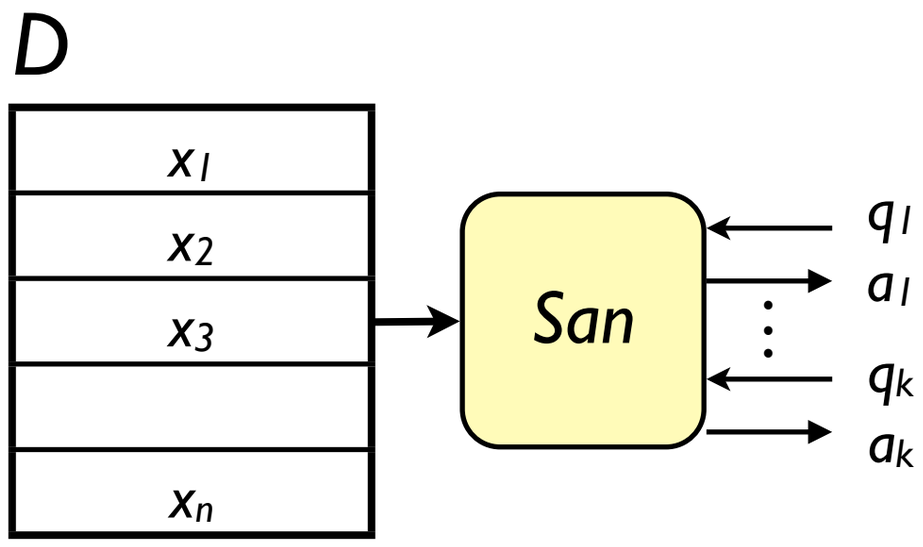

- When is touching the data essential?

Questions for the Workshop

- Managing the privacy budget
  - Composition properties of differential privacy allows us to "divvy up ε" for different uses. But who gets ε? How fast does privacy degrade in reality?
  - We know that answering many "harmless-looking" queries with too little noise can lead to serious attacks [DN03,...]

Thanks! Questions?