Differential Privacy

CPSC 457/557, Fall 13

10/31/13

Hushiyang Liu

Motivation: Utility vs. Privacy

- Era of big data
  - large-size database
  - automatized data analysis
- Utility
  - *"analyze and extract knowledge from data"*
- Privacy
  - sensitive databases, e.g., census, medical, educational, financial, web traffic, OTC drug purchases, query logs, social networking etc.
- Achieve utility while maintain privacy
  - possible?
  - how?

Source: Dwork2011

Motivation: Assumption and Definition

- Analyze data in a privacy-preserving manner
  - assumption: resolved other threats
    - theft, phishing, viruses, cryptanalysis, changing privacy policies …
  - definition of “privacy-preserving” ?

Source: Dwork2011

Motivation: Anonymization?

- "anonymized” or "de-identified”
  - clean off data that is directly linkable to identities
  - non-interactive method
  - vague definition but very broad potential impact (if achieved)

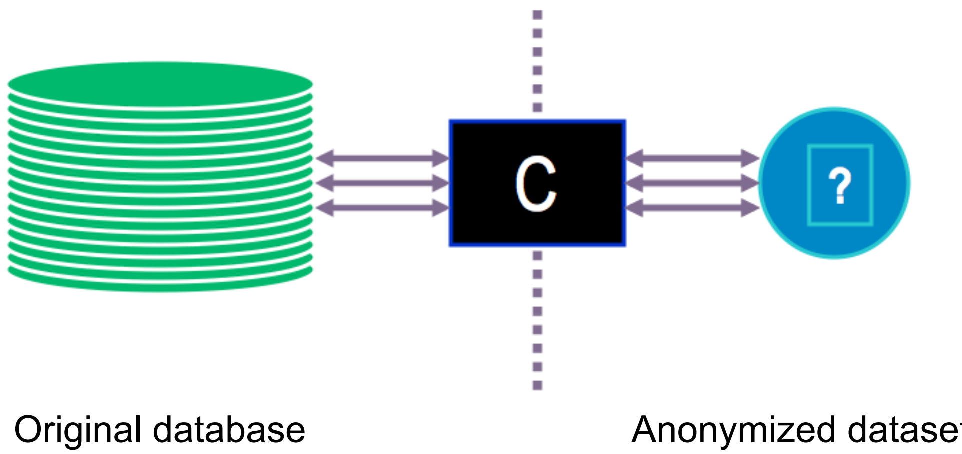

Source: Dwork2011

Motivation: Failure of Anonymization

- Attack against Randomized IDs
  - AOL search data leak of an old woman in Georgia (New York Times, 2006)
    - searcher No. 4417749
    - “numb fingers”
    - “60 single men”
    - “dog that urinates on everything.”
    - “landscapers in Lilburn, Ga”
    - …

- Thelma Arnold
- a 62-year-old widow
- frequently researches medical ailments
- loves her three dogs
- lives in Lilburn, Ga.

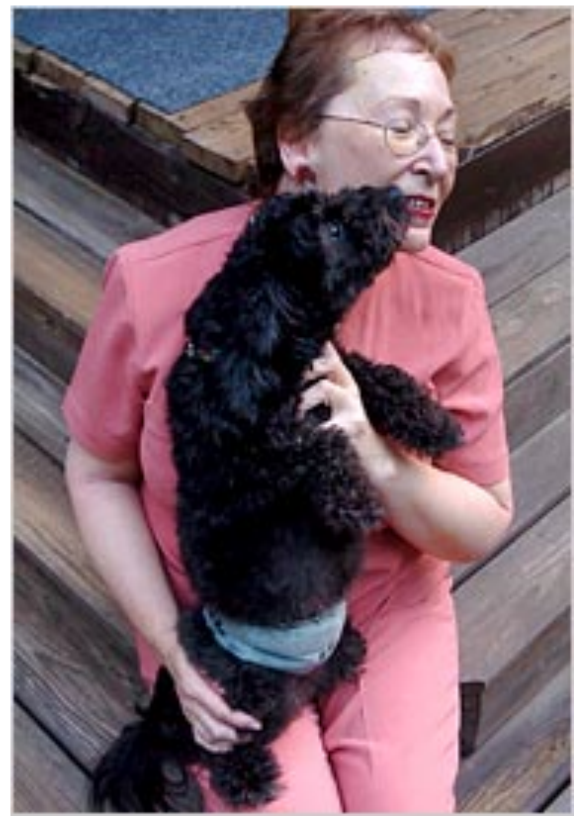

Source: Dwork2011, BZH2006

Motivation: Failure of Anonymization

- Linkage attack: cross-referencing with auxiliary information
  - Massachusetts Governor's medical record – linked "anonymized" HMO data to voter registration data (Latanya Sweeney, 1997)

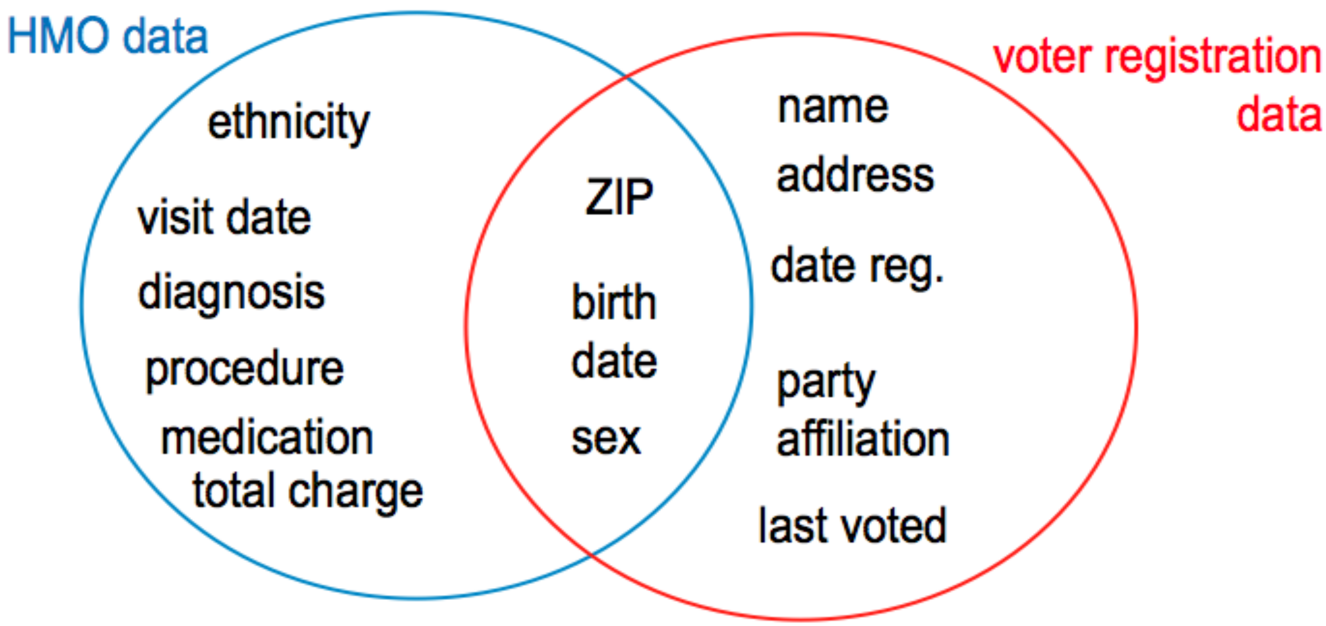

Source: Dwork2011, BZH2006

Motivation: Definitional Failures

- Failure to define privacy
  - failure to account for auxiliary information
  - syntactic and ad hoc
- Need a semantic and "ad omnia" definition that composes automatically and obliviously with (past and future) information

Source: Dwork2011, Dalenius1977, Dwork2006

Motivation: Dalenius's Ad Omnia Guarantee

- Dalenius's Ad Omnia Guarantee [Dalenius1977]
  - *"Anything that can be learned about a respondent from the statistical database can be learned without access to the database."*
  - prior and posterior views about an individual shouldn't change too much
- Provably unachievable [Dwork2006]
  - deductive results
    - "smoking causes cancer" (utility of a database)
    - "Jim smokes" (auxiliary information)
    - "Jim has cancer" (privacy breach!)
  - harm is independent of whether one is in the database

Source: Dwork2011, Dalenius1977, Dwork2006

Motivation: Back to Definitional Failures

- Need a semantic, "ad omnia”, and achievable definition that composes automatically and obliviously with (past and future) information
  - whether or not an analyst interacts with a database => whether or not an individual joins a database
  - differential privacy

Source: Dwork2011, Dwork2006

Differential Privacy

- Definition/Goal: The risk to one's privacy (or in general, any type of risks) should not substantially increase as a result of participating in a statistical database
  - individual privacy
  - privacy budget
  - two "worlds" associated with two databases which differ in only one individual data point (neighboring databases)
- "Differential" refers to the difference between two "worlds"
- Allows for the release of data while meeting a high standard for privacy protection

Source: Dwork2006

Differential Privacy

- Method
  - analyst sends a query to a trusted privacy guard
  - the guard assesses its privacy impact using a special algorithm
  - the guard sends the query to the database and gets back a true answer to that query
  - the guard adds "noise", scaled to the privacy impact, to the answer, and sends the result to the analyst

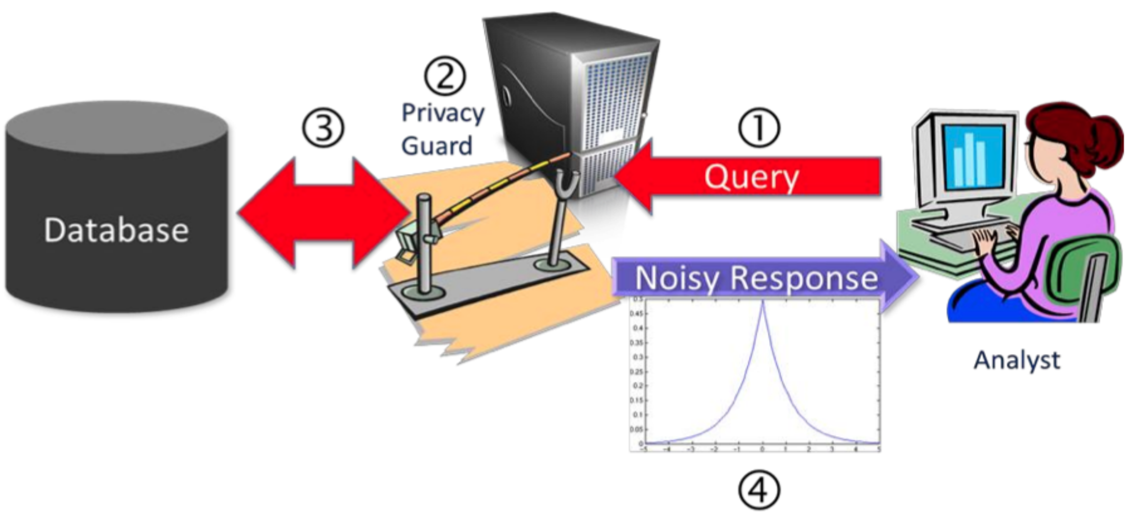

Source: MSFT2013

Algorithm: Basics

- ε-differential privacy for a given result r
  – two neighboring databases D_1 and D_2
  – cannot tell if a result r is from database D_1 or D_2
  – ratio of probabilities should be bounded by e^ε, where ε is a small positive number

$$\frac{P(result = \ r|true\ world = D_1)}{P(result = \ r|true\ world = D_2)} \leq e^{\varepsilon}$$

Source: Dwork2006

Algorithm: Basics

- Global sensitivity Δf
  - f is the *query function* which maps a database to a vector of values (result)

$$\Delta f = \max_{D_1, D_2} \| f(D_1) - f(D_2) \|_1$$

  - Δf is a property of the query function alone
  - sum of the worst-case differences in answers that can be caused by adding or removing one individual from the database
  - a simple example in which the dimension of the result vector is 1
    - f = "how many students scored 100 in the final exam of CS557", D_1 = "all students in CS557", D_2 = "all students in CS557 except Melody"
    - Δf = 1
  - assume that the dimension of the result vector is 1 in the following slides

Source: Dwork2006

Algorithm: Privacy Mechanism

- Add noise to fill the sensitivity gap
  - K_f , a privacy mechanism for a query function f, generates privatized result by computing the real result f(D) and then adding a noise
  - K_f produces a similar distribution of privatized result for two worst-case neighboring databases
  - distributions of possible results from neighboring datasets overlap heavily with each other

$$\frac{P(\mathrm{K}_f(D_1) = r)}{P(\mathrm{K}_f(D_2) = r)} \leq e^{\varepsilon}$$

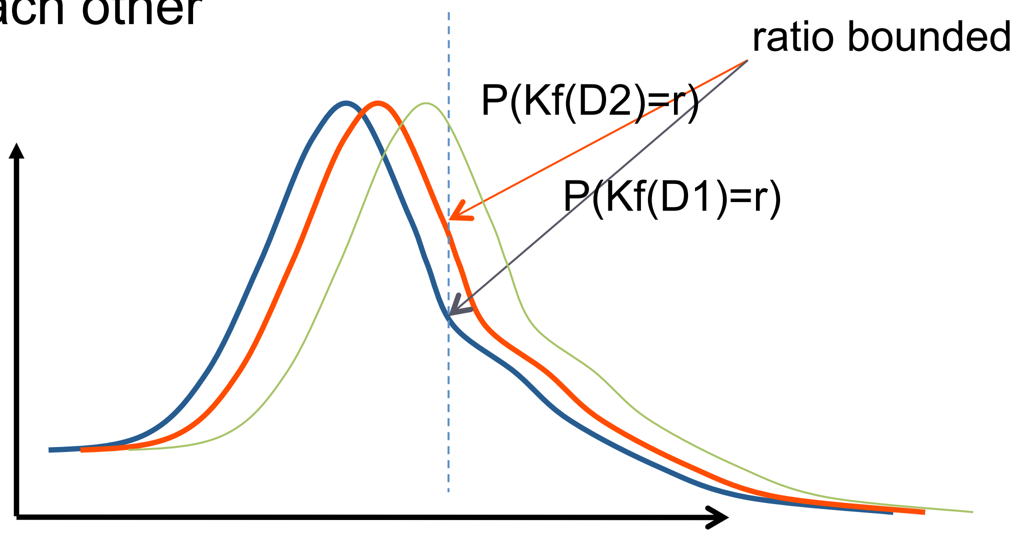

Source: Dwork2006

Algorithm: Choice of Noise

- *Laplacian* noise is an easy way to achieve it
  - *Laplacian* distribution

$$P(x|\mu,\sigma) = \frac{1}{2\sigma} e^{-\frac{|x-\mu|}{\sigma}}$$

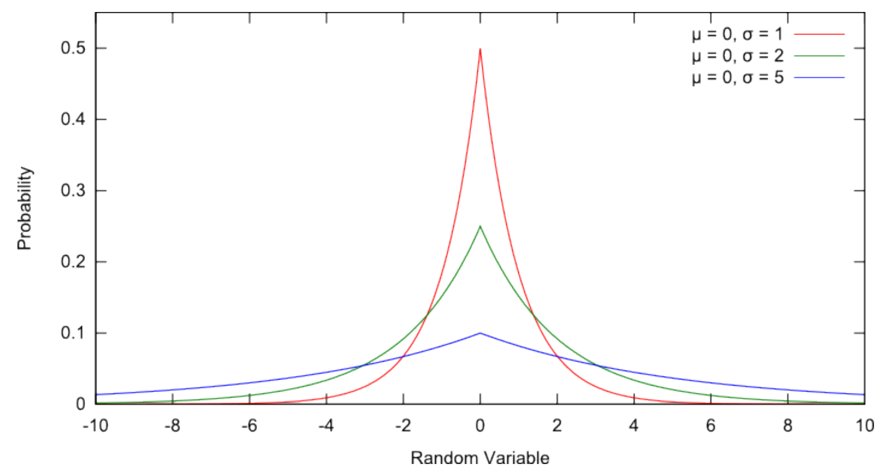

  - privacy mechanism K_f sets

$$\mu = r \text{ and } \sigma = \frac{\varepsilon}{\Delta f}$$

  - K_f produces distribution

$$P(\mathrm{K}_f(D) = r) = \frac{\varepsilon}{2\Delta f} e^{-\frac{|f(D)-r|\varepsilon}{\Delta f}}$$

  - proved in [Dwork2006] that for any pair of neighboring databases D_1, D_2

$$\frac{P(\mathrm{K}_f(D_1) = r)}{P(\mathrm{K}_f(D_2) = r)} \leq e^{\varepsilon}$$

Source: Dwork2006, Wikipedia

Algorithm: Privacy Budget

- ε - privacy budget
  - "Privacy is a nonrenewable resource."
  - predefined privacy variance 1/ε
  - smaller ε means higher privacy

$$\frac{P(\mathrm{K}_f(D_1) = r)}{P(\mathrm{K}_f(D_2) = r)} \leq e^{\varepsilon}$$

- Interactive queries
  - a series of k queries asked by the analyst
  - add noise with variance k/ε to each query [Dwork2006]
  - protect against attack by averaging repeated queries

Source: Dwork2006

Algorithm: Many Others For Better Usage

- When noise makes no sense
  - the function f maps databases to strings, strategies, or trees
  - Exponential Mechanism [MT2007]
- Other algorithms to deal with different cases
  - Statistical Interference
  - Contingency Table Release
  - Halfspace Queries
  - ...

Sources: Dwork2008, MT2007

Application

- Low-error high-privacy DP techniques are applied in
  - Binary Decision Trees
  - Network Trace Analysis
  - Click Query Graphs
  - K-Core Clustering
  - Combinatorial Optimization
  - Frequent Itemset Mining
- Programming platform
  - Privacy Integrated Queries (PINQ) [McSherry2009]
- …

Source: Task2012, McSherry2009

Comment: Evolution

- Underlying data in database remains intact
- Distortion is introduced a posteriori
- Keep track of the cumulative privacy cost
- Good abstraction for analysts to use
- Resilience to all auxiliary information

Source: Dwork2011, MSFT2013

Comment: Limitation

- Narrowness of definition of privacy
  - does not guarantee absolute privacy: deductive results
  - does not guarantee privacy of cohesive group
- Tensions between privacy and utility
  - overwhelming noise
- Complexity of queries
  - “the mean of scores”
- …

Source: Task2012

Discussions

- Do you have a "solution" to the problems of "overwhelming noise" or "complex queries" in DP?

- Can you suggest an alternative protection method? One with a broader definition of privacy?

Discussions

- Do you have a "solution" to the problems of "overwhelming noise" or "complex queries" in DP?
  - ask fewer questions, prune off answers by yourself
  - use result from query with lower sensitivity

- Can you suggest an alternative protection method? One with a broader definition of privacy?

Source: NS2008, NS2007, BDK2007, Task2012

References

1. Fayyad, Usama, Gregory Piatetsky-Shapiro, and Padhraic Smyth. "From data mining to knowledge discovery in databases." AI magazine 17.3 (1996): 37. [FPS1996]
2. Dwork, Cynthia. "A firm foundation for private data analysis." Communications of the ACM 54.1 (2011): 86-95. [Dwork2011]
3. Barbaro, Michael, Tom Zeller, and Saul Hansell. "A face is exposed for AOL searcher no. 4417749." New York Times 9.2008 (2006): 8For. [BZH2006]
4. Dalenius, Tore. "Towards a methodology for statistical disclosure control." Statistik Tidskrift 15.429-444 (1977): 2-1. [Dalenius1977]
5. Dwork, Cynthia. "Differential privacy." Automata, languages and programming. Springer Berlin Heidelberg, 2006. 1-12. [Dwork2006]
6. Microsoft Corporation. "Differential Privacy for Everyone." Retrieved by 2013, http://www.microsoft.com/en-us/download/details.aspx?id=35409http://www.microsoft.com/en-us/download/details.aspx?id=35409. [MSFT2013]
7. Dwork, Cynthia. "Differential privacy: A survey of results." Theory and Applications of Models of Computation. Springer Berlin Heidelberg, 2008. 1-19. [Dwork2008]
8. McSherry, Frank, and Kunal Talwar. "Mechanism design via differential privacy."Foundations of Computer Science, 2007. FOCS"07. 48th Annual IEEE Symposium on. IEEE, 2007. [MT2007]
9. McSherry, Frank D. "Privacy integrated queries: an extensible platform for privacy-preserving data analysis." Proceedings of the 2009 ACM SIGMOD International Conference on Management of data. ACM, 2009. [McSherry2009]
10. Task, Christine. A Practical Beginners" Guide to Differential Privacy. CERIAS Seminar. Purdue University, 2012. [Task2012]
11. Narayanan, Arvind, and Vitaly Shmatikov. "Robust de-anonymization of large sparse datasets." Security and Privacy, 2008. SP 2008. IEEE Symposium on. IEEE, 2008. [NS2008]
12. Backstrom, Lars, Cynthia Dwork, and Jon Kleinberg. "Wherefore art thou r3579x?: anonymized social networks, hidden patterns, and structural steganography." Proceedings of the 16th international conference on World Wide Web. ACM, 2007. [BDK2007]
13. Narayanan, Arvind, and Vitaly Shmatikov. "De-anonymizing social networks."Security and Privacy, 2009 30th IEEE Symposium on. IEEE, 2009. [NS2009]

Thank you

Hushiyang Liu
hushiyang.liu@yale.edu