Differential Privacy

Cynthia Dwork

Mamadou H. Diallo

Overview

- Focus
  - Privacy preservation in statistical databases
  - Goal: to enable the user to learn properties of the population as a whole, while protecting the privacy of the individuals in the sample
- Motivating problem
  - How to reveal useful information about the underlying population, as represented by the database, while preserving the privacy of individuals
  - Previous techniques too powerful
- Approach
  - First define privacy goals, then explore utility
  - Prove the impossibility result
  - Define differential privacy
  - Relate differential privacy to some previous work

Private Data Analysis: The Setting

- Privacy mechanisms models
  - Interactive
    - Data collector is trusted
    - Data collector publishes sanitized data
    - Sanitization techniques:
      data perturbation, sub-sampling, removing identifiers
  - Non-interactive
    - Data collector is trusted
    - Data collector provides an interface
    - Users pose queries about the data
    - Users get noisy data
  - State
    - Powerful results for the interactive approach
    - But, less results for the non-interactive approach

Impossibility of Absolute Disclosure Prevention

- Dalenious privacy definition:
  - Access to a statistical database should not enable one to learn anything about an individual that could not be learned without access
- Semantic Security for cryptosystems (ciphertext indistinguishability)
  - An adversary cannot distinguish pairs of ciphertexts based on the message they encrypt (chosen plaintext attack)
- Observation
  - Semantic security for cryptosystems can be achieved
  - But, Dalenious goal, formalized as a relaxed version of semantic security, cannot be achieved
  - Obstacle: auxiliary information
- Example
  - Sensitive information: exact height
  - Database: average height of women of different nationalities
  - Adversary: access to the DB + auxiliary information
  - Auxiliary information: Terry Gross is two inches shorter than the average Lithuanian woman
- Different between the two: utility requirement

Impossibility of Absolute Disclosure Prevention

- Settings
  - Utility vector: w – (binary vector with k length, answers of questions)
  - Privacy breach
    - Turing machine C
    - **Input**: a description of a distribution D, a database DB, and a string s
    - **Output**: 1 bit,
    - **Adversary Wins**: C(D,DB, s) accepts
  - Auxiliary information generator
    - Turing machine
    - **Input**: D, DB
    - **Output**: z (auxiliary information)
  - Adversary
    - Gets z
    - Has access to DB via the privacy mechanism
    - Modeled as communicating Turing machine
  - Simulator
    - gets z
    - No access to DB
  - Privacy Mechanism: Sam()

Impossibility of Absolute Disclosure Prevention

- Theorem:

  Fix any privacy mechanism San() and privacy breach decider C. There is an auxiliary information generator X and an adversary A such that for all distributions D satisfying Assumption 3 and for all adversary simulators A,*
  Pr[A(D, San(D, DB),X(D, DB)) wins] – Pr[A(D, X(D, DB)) wins] ≥ Δ*
  where Δ is a suitably chosen (large) constant
  The probability spaces: over choices of DB and coin flips of San, X, A, A.

  - **Assumption 3**: For some l satisfying Assumption 2(2b), for any privacy breach y in $\{0, 1\}^l$, the min-entropy of (San(W)|y) is at least k+l, where k is the length of the public strings p produced by the fuzzy extractor.
    w in M and W is any distribution on M

  - **Assuption 2** : There exists an l such that
    2. b) Every DB in D has a privacy breach of length l.

  - Techniques: min-entropy, fuzzy extractors, Hamming distance

Different Privacy

- From absolute to relative guarantees about disclosures
- Differential privacy

  The risk to one's privacy should not substantially increase as a result of participating in a statistical database
- Definition
  - A randomized function K gives ε-differential privacy if for all data sets D1 and D2 differing on at most one element, and all S Range(K),
    Pr[K(D1) in S] ≤ exp(ε) x Pr[K(D2) in S]
  - Observation
    Pr[K(D1) in S] / Pr[K(D2) in S] ≤ exp(ε)
    ln(Pr[K(D1) in S] / Pr[K(D2) in S]) ≤ ε
- Example:
  - The database consulted by an insurance company
  - Should not affect Terry Gross chance of getting insurance
- Definition extension
  - Group privacy
  - c = number of participants
  - Pr[K(D1) in S] / Pr[K(D2) in S] ≤ exp(εc)

Achieving Differential Privacy

- A concrete interactive privacy mechanism achieving ε-differential privacy
  - Query function: f – (simple or complex)
  - Database: X
  - Answer: a = f(X)
- Exponential Noise and the L1-Sensitivity
  - ε-differential privacy achieved by adding a random noise with sensitivity
  - **Sensitivity**:
    The largest change a single participant could have on the output to the query function
    Definition: *for f: D ---> R^d, the L1-sensitivity of f is $\Delta f = max_{D1,D2} \| f(D1) - f(D2) \|_1$ for all D1,D2 differing in at most one element*
    (Techniques work best when Δf is small – least noise)

  - **Density function**:
    Kf: privacy mechanism
    Computes f(X), add noise (scaled symmetric exponential distribution - variance=ρ^2)
    $Pr[K_f(X) = a] = \exp(-\|f(X) - a\|1/\sigma)$
    Implementation: adds symmetric exponential noise to each coordinate of f(X)

  - **Theorem**: for f: D ---> R^d, the mechanism K_f gives (Δf/σ)-differential privacy
  - To achieve ε-differential privacy, choose σ ≥ ε/Δf

Achieving Differential Privacy

- Adaptive adversary
  - fρ: query functions
  - F: deterministic query strategies
  - fρ(X)i: the ith query – (previous responses: ρ1, ρ2, …, ρi-1)
  - F = {f: D ---> (R+)d}
  - Sensitivity of F: $\Delta F = \sup_{\rho} \Delta f_{\rho}$
  - Theorem:
    For query strategy F = {f : D ---> R^d}, the mechanism K_F gives (ΔF/σ)-differential privacy.

Questions?