

Towards Privacy in Public Databases

Shuchi Chawla, Cynthia Dwork,
Frank McSherry, Adam Smith,
Larry Stockmeyer, Hoeteck Wee

Work Done at Microsoft Research

Database Privacy

- Think "Census"
 - Individuals provide information
 - Census Bureau publishes sanitized records
 - Privacy is legally mandated; what utility can we achieve?
- Inherent Privacy vs Utility trade-off
 - One extreme - complete privacy; no information
 - Other extreme - complete information; no privacy
- Goals:
 - Find a middle path
 - preserve macroscopic properties
 - "disguise" individual identifying information
 - Change the nature of discourse
 - Framework for meaningful comparison of techniques

Outline

- Definitions

- privacy, defined in the breach
- sanitization requirements
- utility goals

- Example: Recursive Histogram Sanitizations

- description of technique
- a robust proof of privacy

- Work in Progress

- extensions and impossibility results
- dealing with auxiliary information

Outline

- Definitions
 - privacy, defined in the breach
 - sanitization requirements
 - utility goals
- Example: Recursive Histogram Sanitizations
 - description of technique
 - a robust proof of privacy
- Work in Progress
 - extensions and impossibility results
 - dealing with auxiliary information

What do WE mean by privacy?

- [Ruth Gavison] Protection from being brought to the attention of others
 - inherently valuable
 - attention invites further privacy loss
- Privacy is assured to the extent that one blends in with the crowd
- Appealing definition; can be converted into a precise mathematical statement...

A Geometric View

- Abstraction:

- Database consists of points in high dimensional space \mathbb{R}^d
- Points are unlabeled
you are your collection of attributes
- Distance is everything
points are similar if and only if they are close (L_2 norm)

- Real Database (RDB), private

n unlabeled points in d-dimensional space d
» as number of sensitive attributes


- Sanitized Database (SDB), public

n' new points, possibly in a different space

The Isolator - Intuition

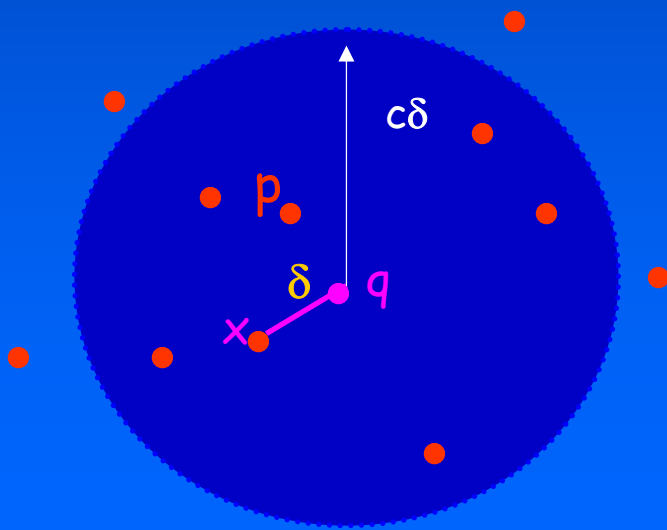
- On input SDB and auxiliary information, adversary outputs a point $q \in \mathbb{R}^d$
- q “isolates” a real DB point x , if it is much closer to x than to x 's near neighbors
 - q fails to isolate x if q looks roughly as much like everyone in x 's neighborhood as it looks like x itself
 - Tightly clustered points have a smaller radius of isolation

RDB



Isolation - the definition

- $I(SDB, aux) = q$
- x is isolated if $B(q, c\delta)$ contains fewer than T other points from RDB



c - privacy parameter; eg, 4

Requirements for the sanitizer

- No way of obtaining privacy if AUX already reveals too much!
- Sanitization procedure compromises privacy if giving the adversary access to the SDB considerably increases its probability of success
- Definition of "considerably" can be forgiving
- Made rigorous by quantification over adversaries, distributions, auxiliary information, sanitizations, samples:
 - $\forall D \forall I \exists I'$ w.h.p. $D \forall aux z$
 $\sum_x |\Pr[I(SDB, z) \text{ isolates } x] - \Pr[I'(z) \text{ isolates } x]|$ is small
 - Provides a framework for describing the power of a sanitization method, and hence for comparisons

Utility Goals

- Natural approaches

- pointwise proofs of specific utilities
 - averages, medians, clusters, regressions,...
- prove there is a large class of interesting tests for which there are good approximation procedures using sanitized data

- Our Results

- concrete pointwise results on histograms and clustering;
- connection to data streaming algorithms that use exponential histograms

Outline

- Definitions

- privacy, defined in the breach
- sanitization requirements
- utility goals

- Example: Recursive Histogram Sanitizations


- description of technique
- a robust proof of privacy

- Work in Progress

- extensions and impossibility results
- dealing with auxiliary information

Recursive Histogram Sanitization

- $U = d$ -dim cube, side = 2
- Cut into 2^d subcubes
 - split along each axis
 - subcube has side = 1
- For each subcube
 - if number of RDB points $> 2T$
 - then recurse
- Output: list of cells and counts



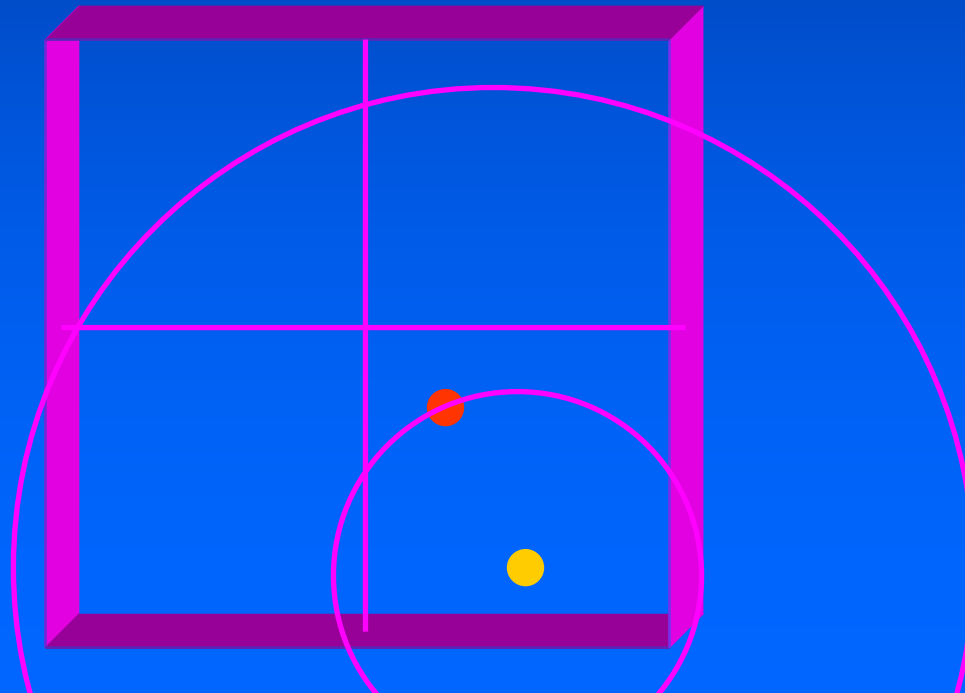
Recursive Histogram Sanitization

- **Theorem:** $\exists c$ s.t. if n points are drawn uniformly from U , then recursive histogram sanitizations are safe with respect to c -isolation:
 $\Pr[\text{I(SDB) succeeds}] \cdot \exp(-d)$.

Safety of Recursive Histogram Sanitization

● Rough Intuition

- Expected distance $||q-x||$ is \approx diameter of cell.
- Distances tightly concentrated around mean.
- Multiplying radius by c captures almost all the parent cell - contains at least $2T$ points.



Proof is Very Robust

- Extends to many interesting cases
 - non-uniform but bounded-ratio density fns
 - isolator knows constant fraction of attributes
 - isolator know lots of RDB points
 - isolation in few attributes
 - weak bounds
- Can be adapted to “round” distributions
 - with effort; Work in Progress [w/ Talwar]

Outline

- Definitions

- privacy, defined in the breach
- sanitization requirements
- utility goals

- Example: Recursive Histogram Sanitizations

- description of technique
- a robust proof of privacy

- Work in Progress

- extensions and impossibility results
- dealing with auxiliary information

Extensions & Impossibility *

- **Relaxed Definition of Isolation**
 - Adversary chooses a small set of attributes on which to isolate; increase c accordingly; histograms still private
- **Impossibility Results**
 - Impossibility of all-purpose sanitizers
 - Interesting utilities that have no privacy-preserving sanitization (cf. SFE)
- **Utility**
 - exploit literature (eg, Indyk+) on power of randomized histograms; extend to histograms for round distributions (how to randomize?)
- **Extensive Work on Round Sanitizations**
 - clustering results
 - privacy via cross-training (done for cubes)

* with assorted collaborators, eg, N,N,S,T

Auxiliary Information

- Protection against isolation yields protection against learning a key for a population unique
 - isolation on a subspace does not imply isolation in the full-dimensional space ...
 - ... but aux may contain other DBs that can be queried to learn remaining attributes
 - definition mandates protection against all possible aux
 - satisfy def) can't learn key

Connection to Real World

- Very hard to provide good sanitization in the presence of arbitrary aux
 - Provably impossible in general
 - Anyway, can probably already isolate people based solely on aux
 - Suggests we need to control aux
- How should we redesign the world?
 - Maybe OK to give data to really trustworthy and audited agency, but what about other entities?

Two Tools

- Secure Function Evaluation [Yao, GMW]
 - Technique permitting Alice, Bob, Carol, and their friends to collaboratively compute a function f of their private inputs $\xi = f(a,b,c,\dots)$.
 - eg, $\xi = \text{sum}(a,b,c, \dots)$
 - Each player learns only what can be deduced from ξ and her own input to f
- SuLQ databases [Dwork, Nissim]
 - Provably preserves privacy of attributes when the rows of the database are mutually independent
 - Powerful [DN; Blum, Dwork, McSherry, Nissim]

Our Data, Ourselves

- Individuals maintain their own data records
 - join a DB by setting an appropriate attribute

0	4	6	3	...	1	0	...
---	---	---	---	-----	---	---	-----

- Statistical queries via a SFE(SuLQ)
 - privacy of SuLQ query) this SFE is "safe"
 - the SFE is just Sum (easy!)
- Individuals ensure
 - data take part in sufficiently few queries
 - sufficient random noise is added



Summary

- Definitions

- defined isolation and sanitization

- Recursive Histogram Sanitizations

- described approach and sketched a robust proof of privacy for a special distribution
- Proof exploits high dimensionality (# columns)

- Additional results

- sanitization by perturbation, impossibility results, utility via data streaming algorithms

- Setting the Real World Context

- discussed a radical view of how data might be organized to prevent a powerful class of attacks based on auxiliary data
- SuLQ tool exploits large membership (# rows)