LA-UR- 09- 07903

Approved for public release;
distribution is unlimited.

Title:	What Can We Learn Privately?
Author(s):	Shiva Kasiviswanthan: CCS-#, Z# 209013, LANL Amos Beimel: Ben-Gurion University Homin Lee: Columbia University Kobbi Nissim: Ben-Gurion University Sofya Raskhodnikova: Pennsylvania State University Adam Smith: Pennsylvania State University
Intended for:	To be Presented: Institute of Mathematical Sciences

Los Alamos
NATIONAL LABORATORY
EST. 1943

Form 836 (7/06)

What Can We Learn Privately?

Shiva Kasiviswanathan
Los Alamos National Lab

Covers Joint work with:
Amos Beimel (Ben-Gurion)
Homin Lee (Columbia)
Kobbi Nissim (Ben-Gurion)
Sofya Raskhodnikova (Penn State)
Adam Smith (Penn State)

Database Privacy: The Setting

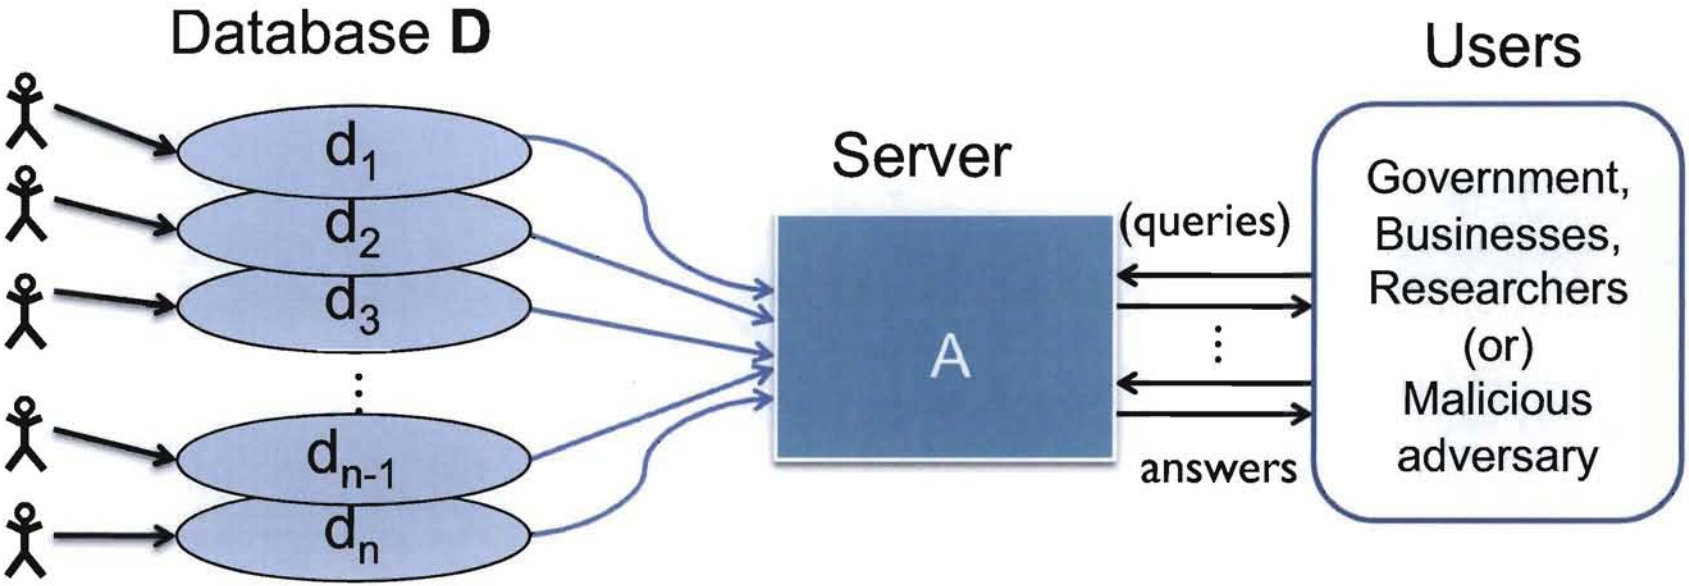

- Database **D** = (d_1,d_2, …,d_n) (a table of n rows)
- Each element is from some domain X
- X can be numbers, categories, tax forms, etc

Privacy in Statistical Databases

Typical examples:

- Census
- Civic archives
- Medical records
- Search information
- Communication logs
- Social networks
- Genetic databases
- …

Benefits: **Social & Economic**

Privacy: **Fundamental right**

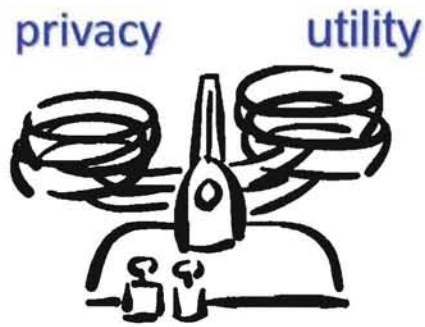

Example 1: Google vs. Viacom

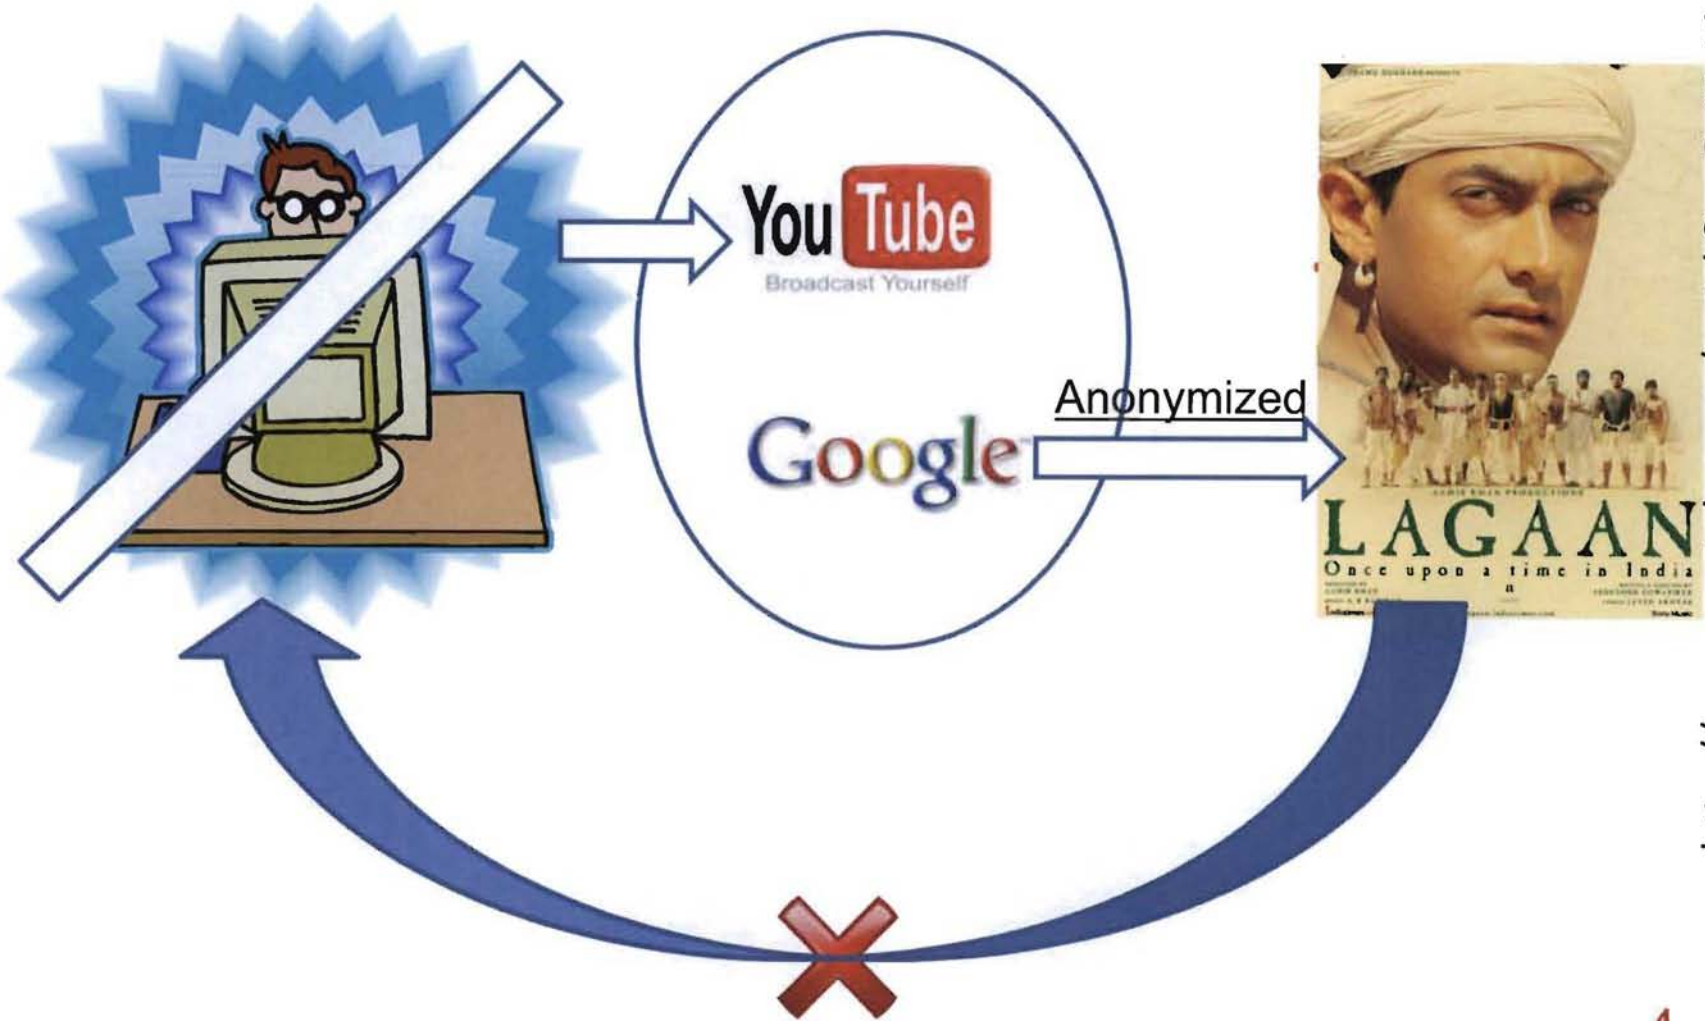

http://www.youtube.com/watch?v=rBsCJpVkaEE

Example 2: AOL Disaster

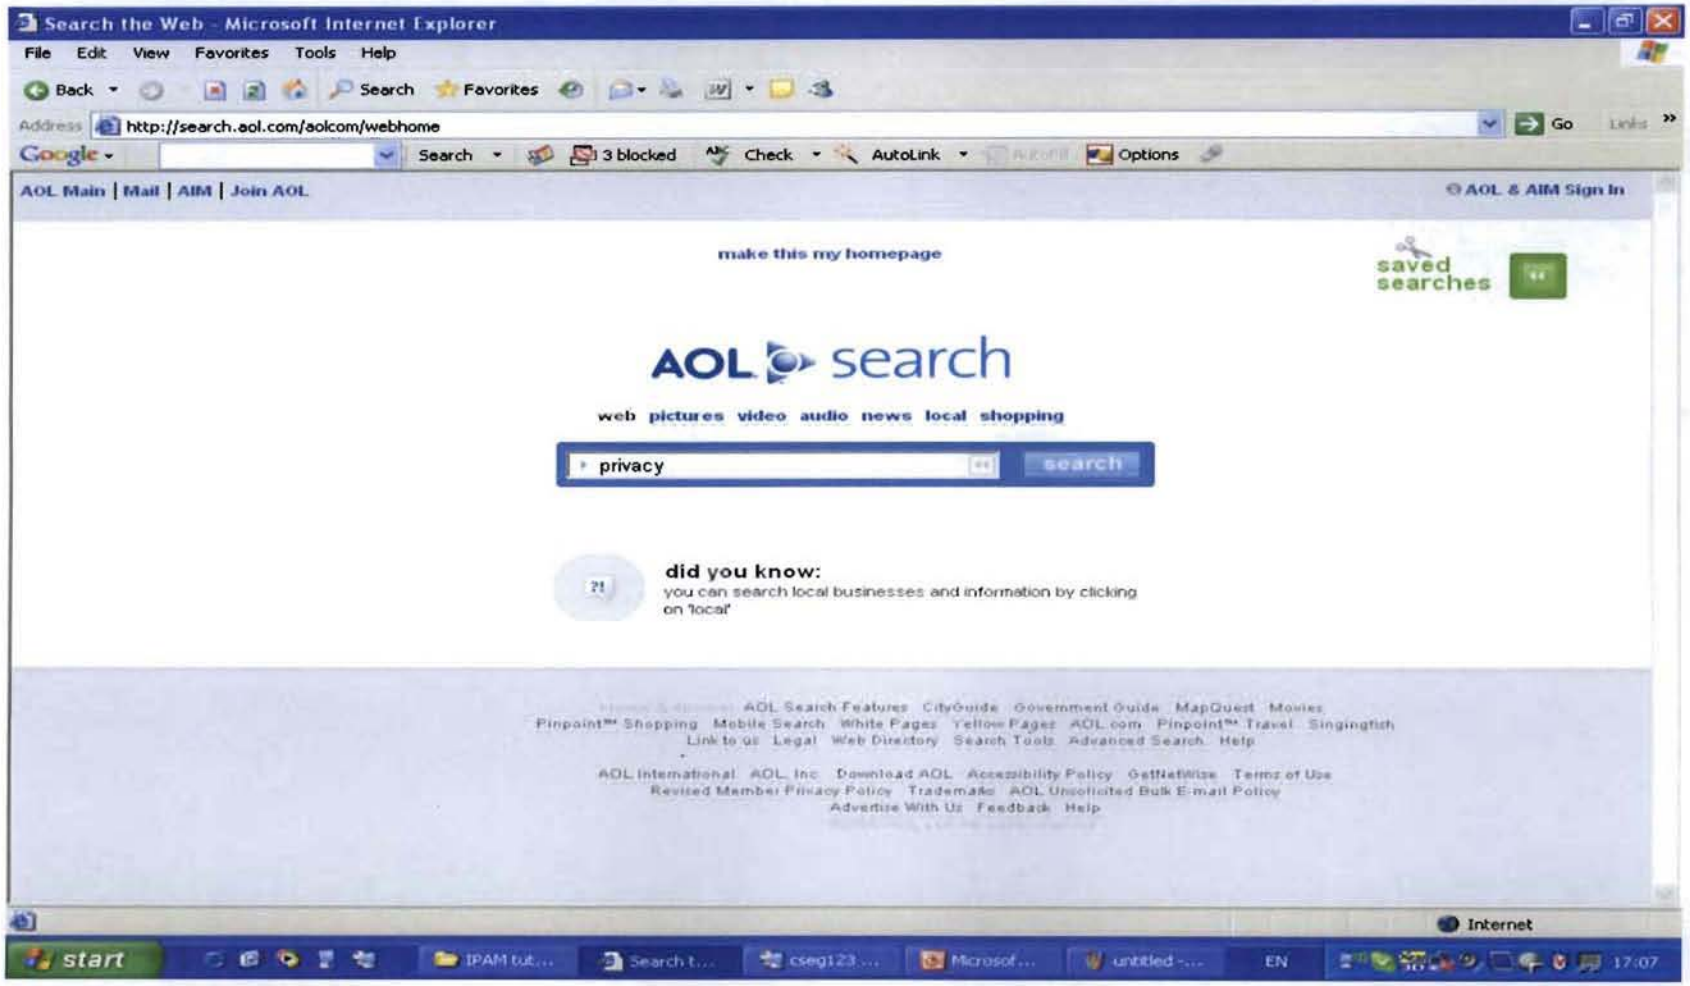

AOL Search History Release

- 650,000 users, 20 Million queries, 3 months
- Goal: provide query log data that is based on real users
- Privacy: Identifying information replaced with random id's

4417749	best dog for older owner	3/6/2006 11:48:24	1	http://www.canismajor.com
4417749	best dog for older owner	3/6/2006 11:48:24	5	http://dogs.about.com
4417749	landscapers in lilburn ga.	3/6/2006 18:37:26		
4417749		effects of nicotine 3/7/2006		
19:17:19		6 http://www.nida.nih.gov		
4417749	best retirement in the world	3/9/2006 21:47:26	4	http://www.escapeartist.com
4417749	best retirement place in usa	3/9/2006 21:49:37	10	http://www.clubmarena.com
4417749	nicotine effects on the body	3/26/2006 10:31:15	2	http://health.howstuffworks.com
4417749	wrinkling of the skin	3/26/2006 10:38:23		
4417749	mini strokes	3/26/2006 14:56:56	1	http://www.ninds.nih.gov
4417749	panic disorders	3/26/2006 14:58:25		
4417749	jarrett t. arnold eugene oregon	3/23/2006 21:48:01	2	http://www2.eugeneweekly.com
4417749	jarrett t. arnold eugene oregon	3/23/2006 21:48:01	3	http://www2.eugeneweekly.com
4417749	plastic surgeons in gwinnett county	3/28/2006 15:04:23	1	http://www.wedalert.com
4417749	plastic surgeons in gwinnett county	3/28/2006 15:31:00		
4417749	60 single men	3/29/2006 20:11:52	6	http://www.adultlovecompass.com
4417749	60 single men	3/29/2006 20:14:14		
4417749	clothes for 60 plus age	4/19/2006 12:44:03		
4417749	clothes for age 60	4/19/2006 12:44:41	10	http://www.news.cornell.edu
4417749	clothes for age 60	4/19/2006 12:45:41		
4417749	lactose intolerant	4/21/2006 20:53:51	2	http://digestive.niddk.nih.gov
4417749	lactose intolerant	4/21/2006 20:53:51	10	http://www.netdoctor.co.uk
4417749	dog who urinate on everything	4/28/2006 13:24:07	6	http://www.dogdaysusa.com
4417749	fingers going numb	5/2/2006 17:35:47		

A Face Is Exposed for AOL Searcher No. 4417749

By MICHAEL BARBARO and TOM ZELLER Jr.
Published: August 9, 2006

The New York Times

Buried in a list of 20 million Web search queries collected by AOL and recently released on the Internet is user No. 4417749. The number was assigned by the company to protect the searcher's anonymity, but it was not much of a shield.

No. 4417749 conducted hundreds of searches over a three-month period on topics ranging from "numb fingers" to "60 single men" to "dog that urinates on

Name: Thelma Arnold
Age: 62
Widow
Residence: Lilburn, GA

Lines of Work on Database Privacy

Huge amount of work in "traditional" fields

- In statistics (statistical disclosure control)
- In data mining (privacy-preserving data mining)
- Problem: no precise privacy definition

In cryptography (private data analysis)
[DiNi, DN, BDMN, DKMMN, CDMT, DMNS, Dw, BCDKMT, MT, NRS, MKGAV, BLR,.....]

Goal: Rigorous privacy guarantees

Definition: Differential Privacy [DMNS06]

A randomized algorithm A is ε-differentially private if

- for all databases **D**, **D'** that differ in one element
- for all sets of answers S

$$\Pr[A(\mathbf{D}) \in S] \leq e^{\varepsilon} \Pr[A(\mathbf{D'}) \in S]$$

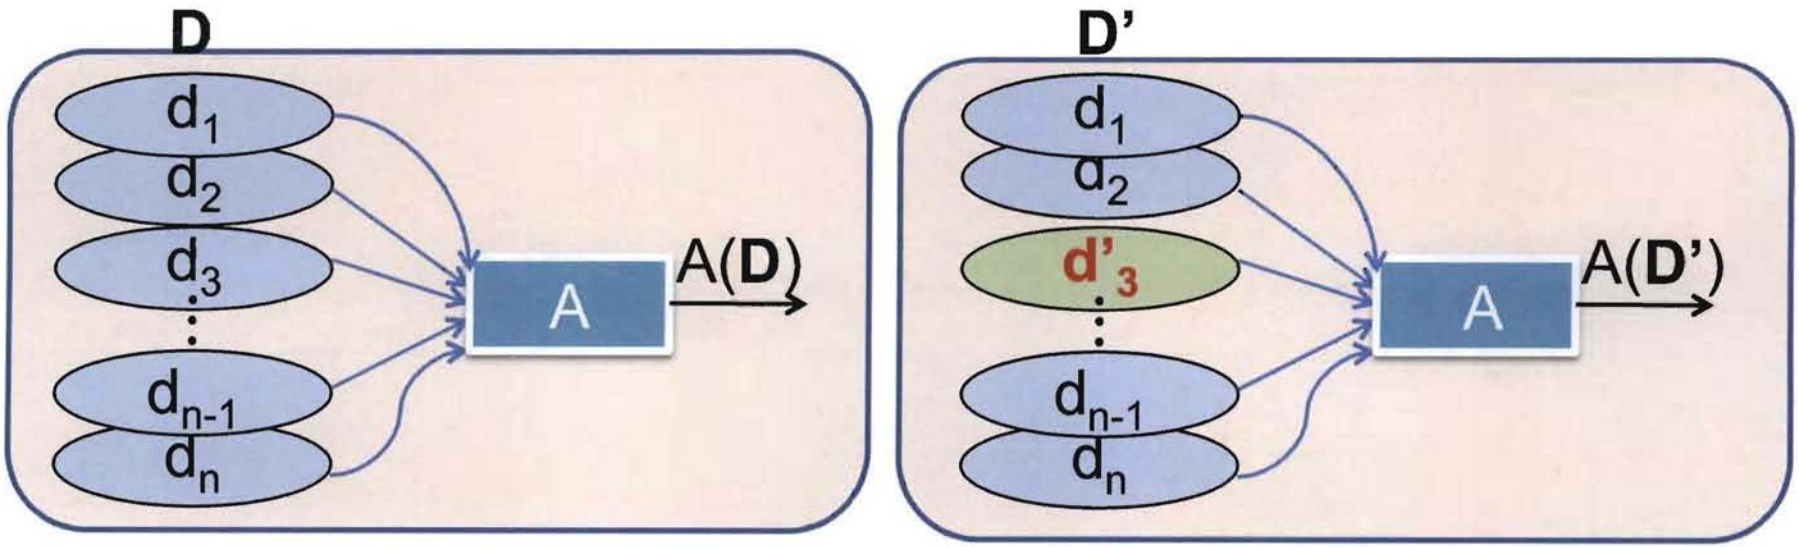

Definition: Differential Privacy [DMNS06]

- ε cannot be too small (think 1/10, not $1/10^5$)
- Distance measure on distributions matters
- This is a condition on the **algorithm** (process) A

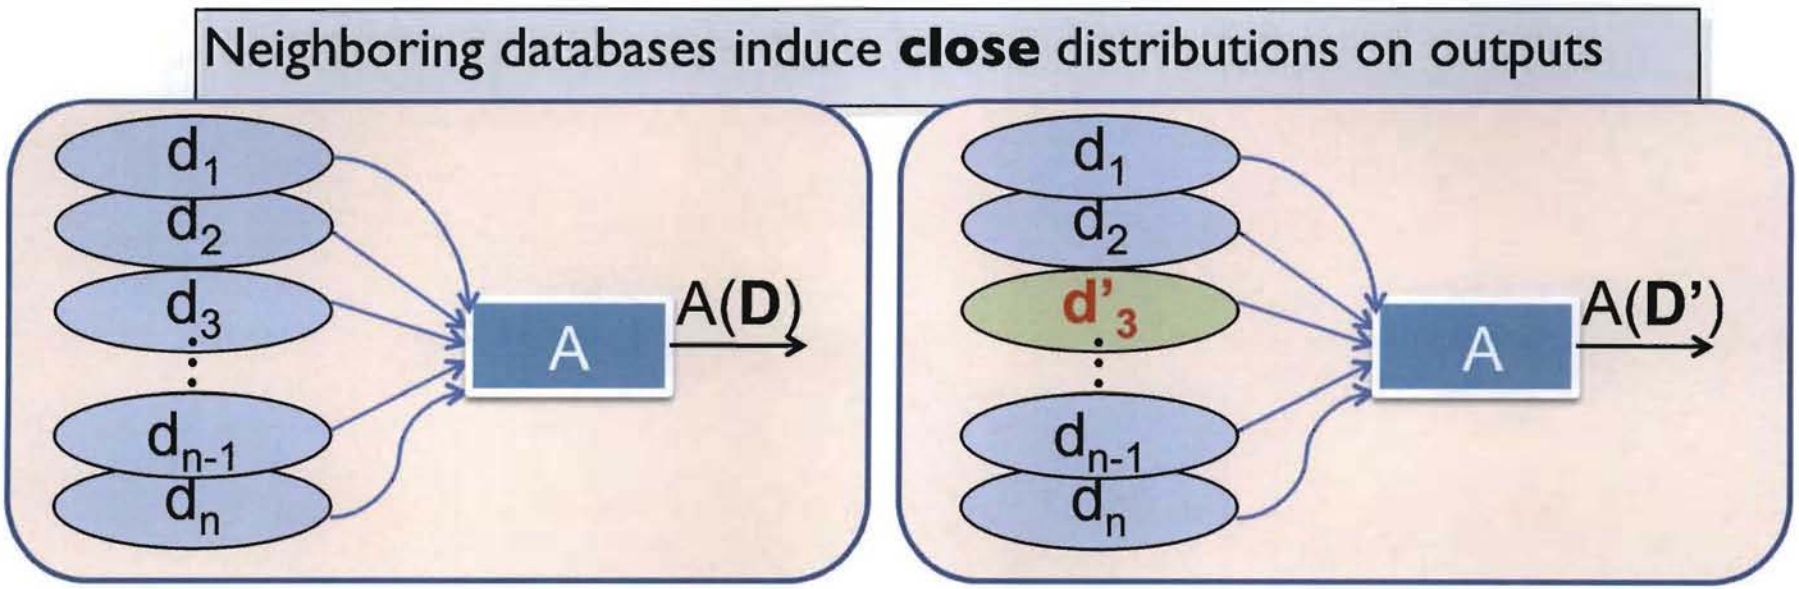

Why is this a Good Definition?

~~Ideal: Your beliefs about me are the same after you see the output as they were before~~

- Suppose you know I am as tall as average Indian
  - You could learn my height from database!
  - But it didn't matter whether or not my data was part of it.

Theorem (DN): Prefect privacy is unachievable

Differential privacy implies

You learn the same things about me whether or not I am in the database

Why is this a Good Definition?

- Consider an intruder trying to infer personal information
  - "Background knowledge" = prior distribution on data **D**
  - "Conclusions you draw" = posterior $p(\cdot|output)$
  - Experiment 0: Run A(**D**)
  - Experiment *i*: Run A($\mathbf{D}_{-i}$) where $\mathbf{D}_{-i} = (d_1,..,d_{i-1},*,d_{+1},.,d_n)$

Lemma: $\forall$ prior, $\forall$ output, $p_0(\cdot|output) \approx p_i(\cdot|output)$

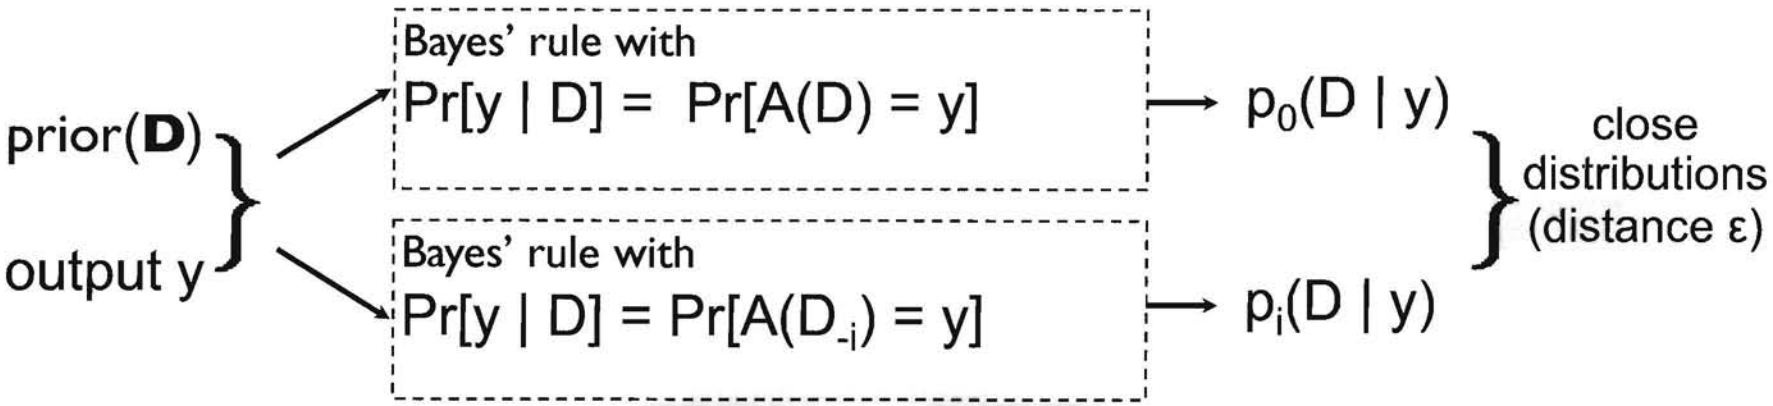

What Can We Compute Privately?

Exception

Focused on function evaluation: compute f(**D**)

[DiNi03,BDMN05,DMNS06,Dw06,NRS07,DMT07,BCDKMT07,MT07,BLR08]

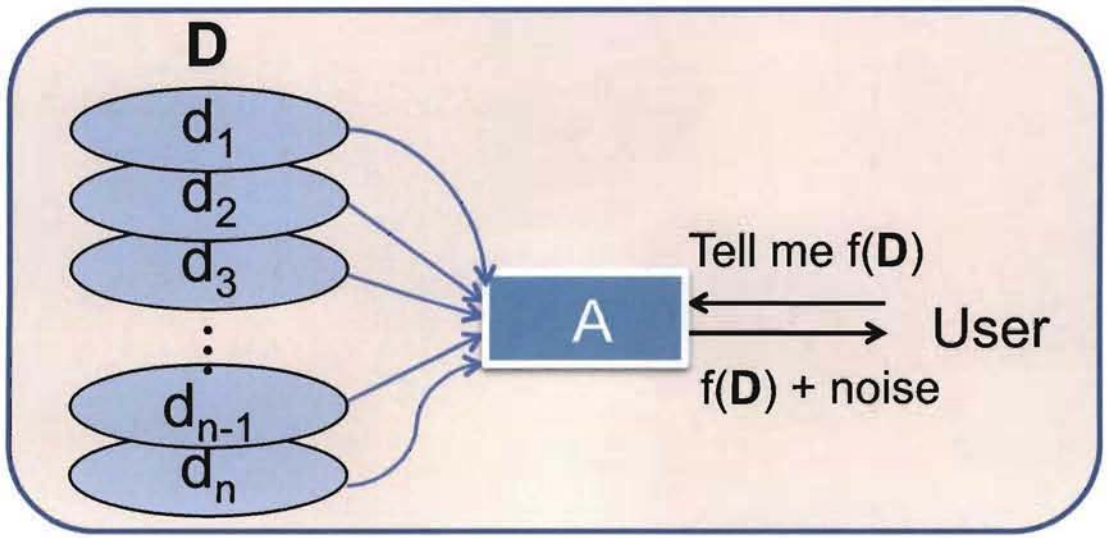

Starting Point: Study which other computational tasks can be performed in a privacy preserving manner

This Talk

"Differential" privacy

- Handles arbitrary external information
- What can we compute privately?

Learning algorithms that respect privacy

Private Learning

- Goal: machine learning algorithms that protect the privacy of individual examples (people, organizations,..)

Private Learning
- Utility: Learning guarantees
- Privacy: Differential privacy guarantees

This work:

✧Characterize classification problems learnable privately

✧Understand power of popular privacy settings for learning

Private Learning

- Goal: machine learning algorithms that protect the privacy of individual examples (people, organizations,..)

Private Learning
- Utility: Learning guarantees
- Privacy: Differential privacy guarantees

Focus of this talk:

✧Characterize classification problems learnable privately

✧Understand power of popular privacy settings for learning

Our Results 1: What is Learnable Privately

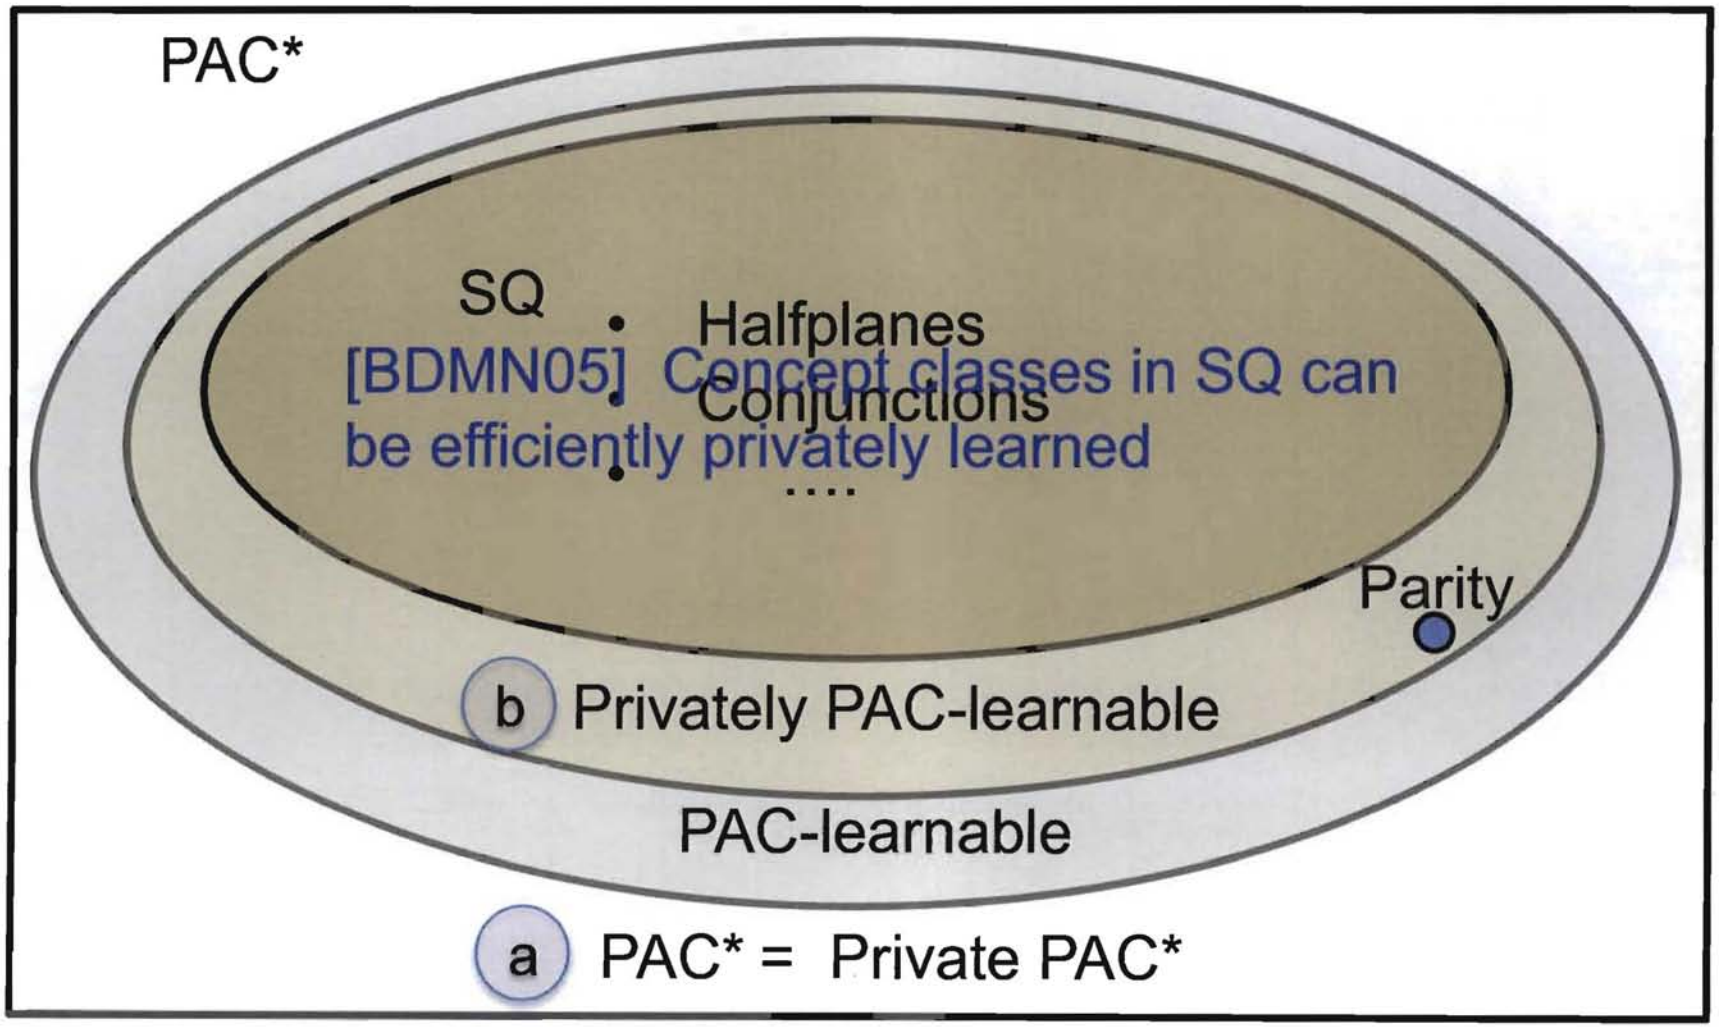

Learning and Private Learning

Learning Setting: Classification*

- Bank needs to decide which applicants are bad credit risks
- Goal: given sample of labeled data (past customers), produce good prediction rule ("hypothesis") for future loan applicants

% down	High Debt	Other accts	Mmp/inc	Good Risk?
10	No	Yes	0.32	Yes
10	No	No	0.25	Yes
5	Yes	No	0.30	No
20	No	Yes	0.31	Yes
10	No	No	0.25	Yes

example y_i

label z_i

- Reasonable rules given this data:
  - Predict YES iff 100*(Mmp/inc) – (% down) < 25
  - Predict YES iff (!High Debt) AND (% down > 5)

* Example taken from Avrim Blum, FOCS03 tutorial

PAC (Probabilistically Approximately Correct) Model
[Valiant84]

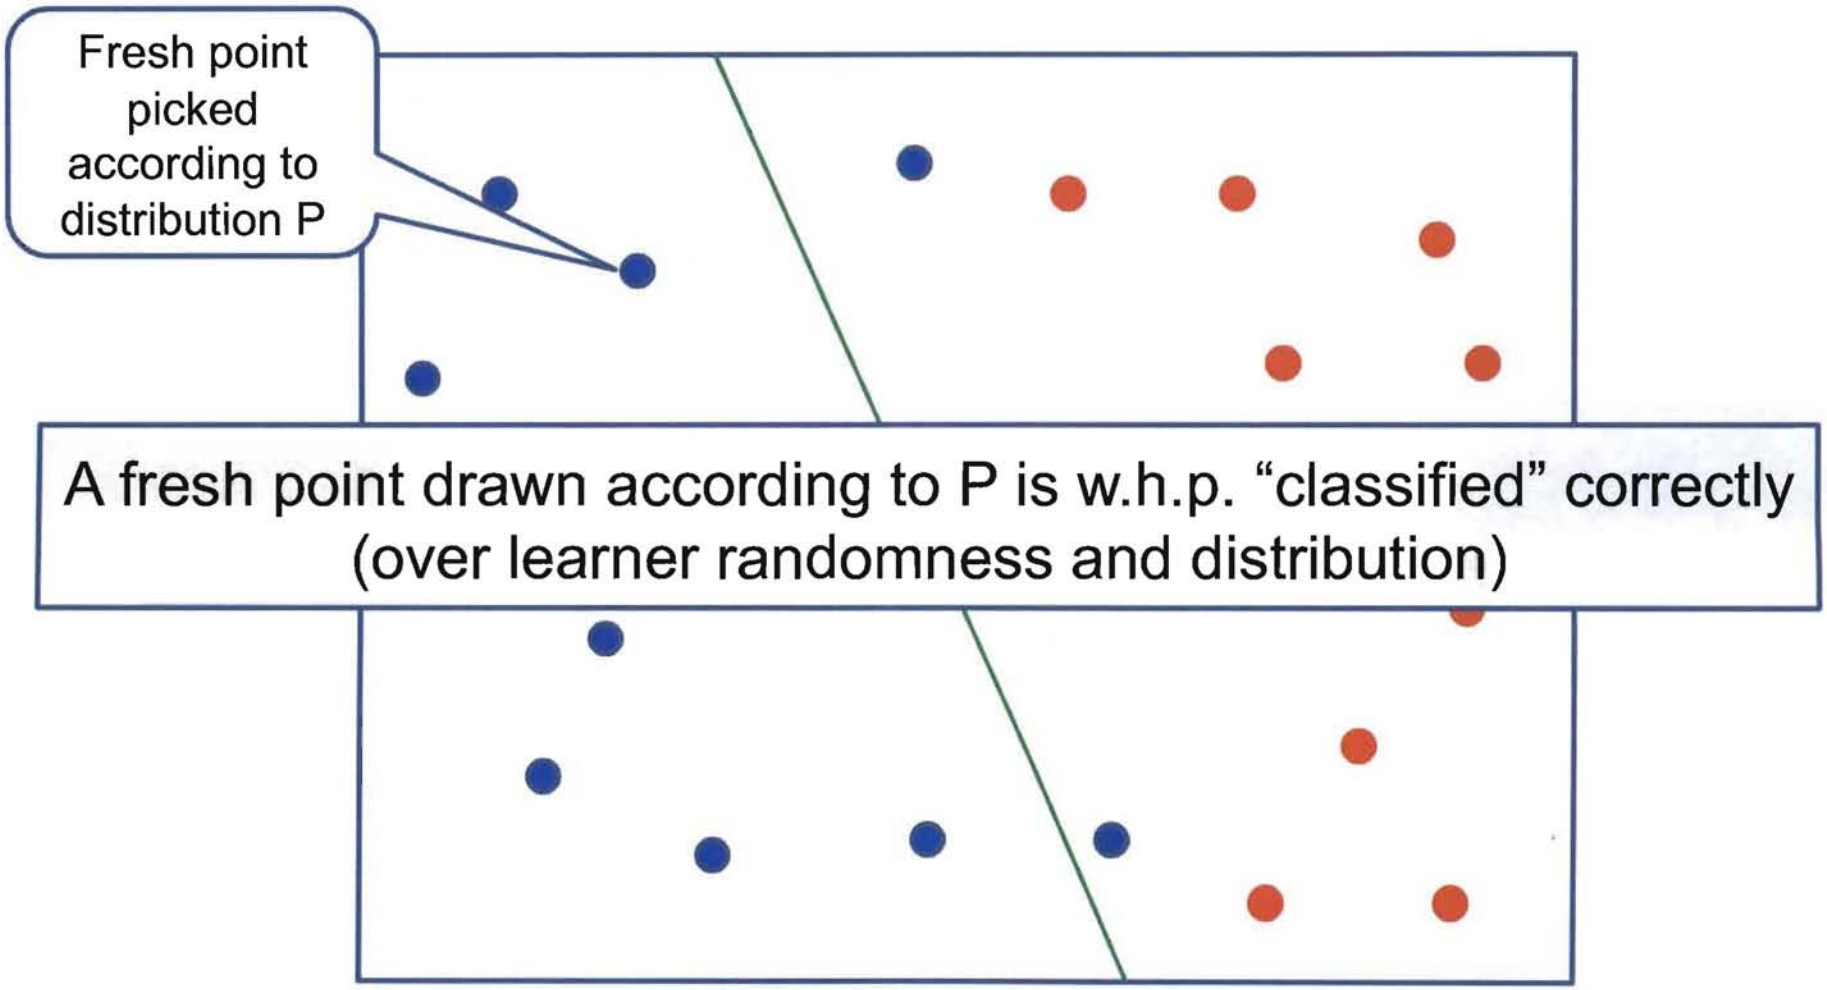

PAC Learning Definition

Given distribution P over examples, labeled by function c,
hypothesis h is good if it mostly agrees with c

$Pr_{y \sim P}[h(y) = c(y)]$ is close to 1 w.h.p.

Let C be a set of concepts $\{c : X \rightarrow \{0,1\}\}$

Algorithm A (**PAC***) learns C if, for every c in C

- given poly examples drawn from P, labeled by c

  of polynomial size
- A outputs a good hypothesis $\hat{}$ w.h.p. ~~in poly time~~

Private Learning

Input: Database: **D** = $(d_1, d_2, \ldots, d_n)$ where

$d_i = (y_i, z_i)$, where $y_i \sim P$, $z_i = c(y_i)$ (z_i is the label on example y_i)

% down	High Debt	Other accts	Mmp /inc	Good Risk?
10	No	Yes	0.32	Yes
10	No	No	0.25	Yes
25	No	No	0.30	Yes
20	No	Yes	0.31	Yes
10	No	No	0.25	Yes

Output is a hypothesis e.g.

Predict Yes if 100*(Mmp/inc)-(% down) <25

Algorithm A privately PAC* learns concept class C if:

- Utility: Algorithm A PAC* learns concept class C

Distributional guarantee

- Privacy: Algorithm A is ε-differentially private

Worst-case guarantee

PAC* = Private PAC*

Theorem: Every PAC* learnable concept class can be learned privately, using a polynomial number of samples

Proof: Adapt exponential mechanism [MT07]:

$q(\mathbf{D},h)$ = # of examples in **D** correctly classified by hypothesis h

Algorithm A: Output hypothesis h from concept class C with probability~ $\mathbf{e}^{\varepsilon \cdot q(\mathbf{D},h)}$

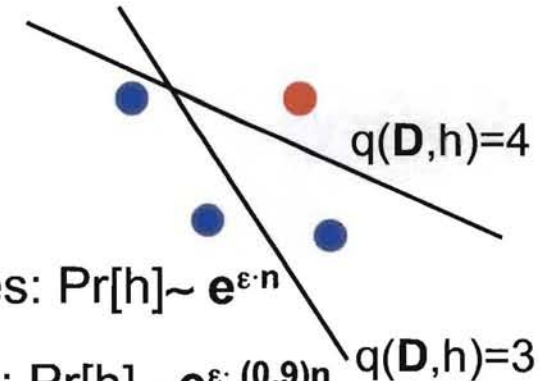

Privacy: for any hypothesis h,

- Good hypotheses correctly label all examples: Pr[h]~ $\mathbf{e}^{\varepsilon \cdot n}$
- Bad hypotheses mislabel ≥10% of examples: Pr[h]~ $\mathbf{e}^{\varepsilon \cdot (0.9)n}$

$$\frac{\Pr[h \text{ is output on input } \mathbf{D}]}{\Pr[h \text{ is output on input } \mathbf{D}']} \leq e^{\varepsilon \cdot q(\mathbf{D},h)} \leq e^{\varepsilon \cdot q(\mathbf{D}',h)} \leq e^{\varepsilon} \text{ polynomial}$$

~~Sufficient to ensure $n >> \log(\#\text{bad hypotheses})$~~

≤ output length of non-private learner

Theorem: Every PAC* learnable concept class can be learned privately, using a polynomial number of samples

Some Observations:

- Above theorem is a private version of Occam's razor
- Can be extended to VC dimension analogue [BLR08]

Power of Popular Privacy Settings

Basic Privacy Settings

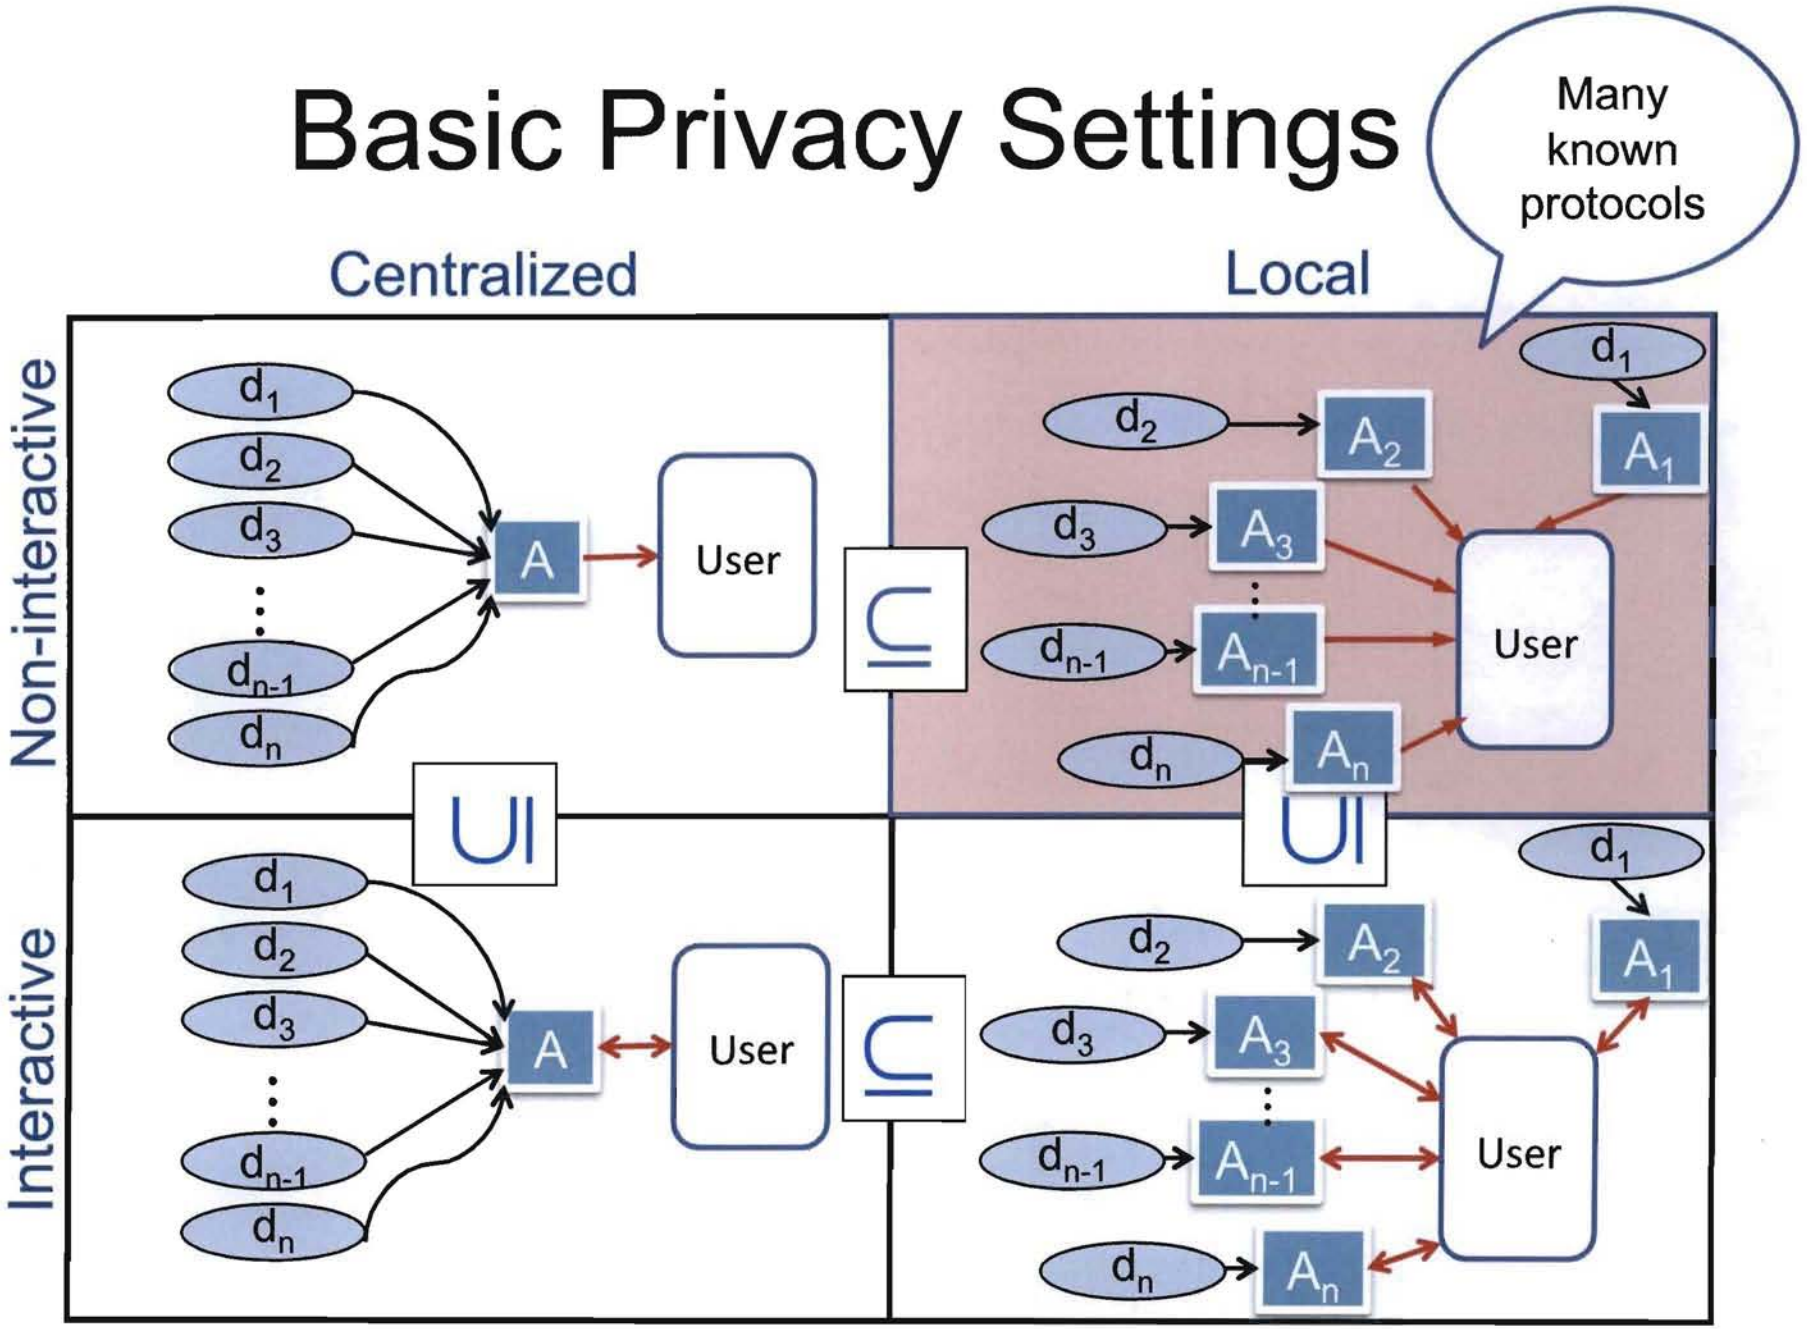

Our Results 2: Local Model

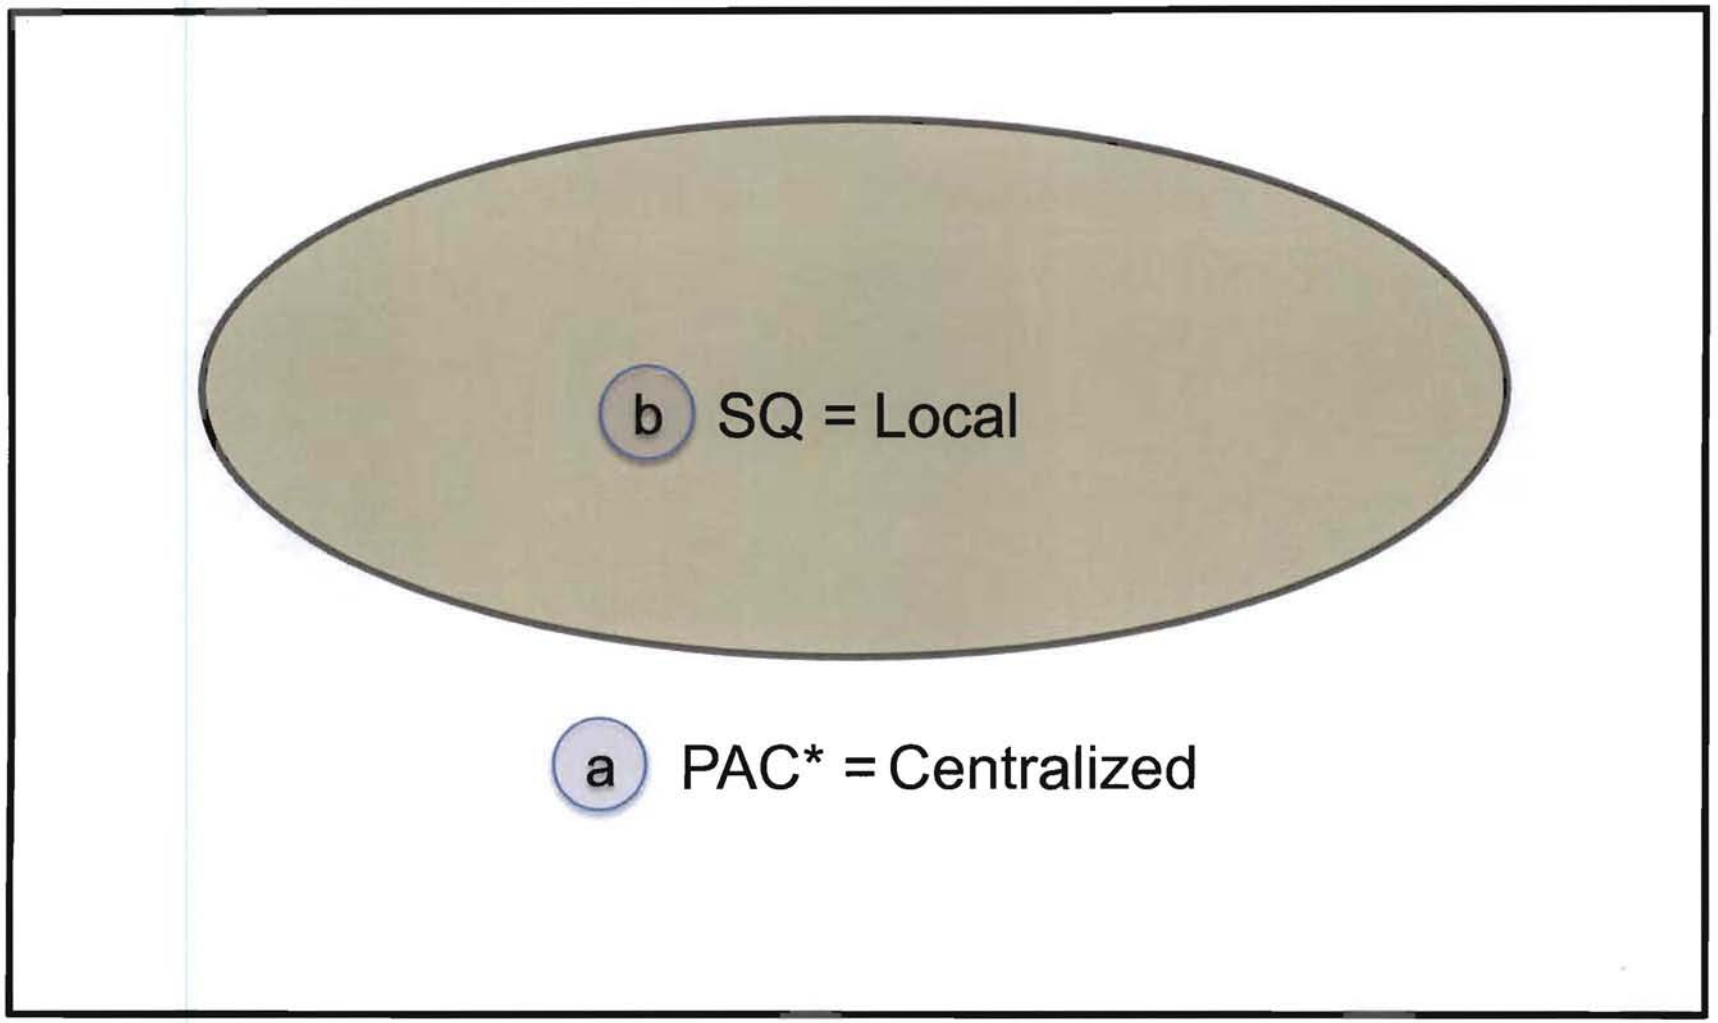

Statistical Query (SQ) Model [Kearns93]

Same guarantees as PAC model, but algorithm no longer has access to individual labeled examples

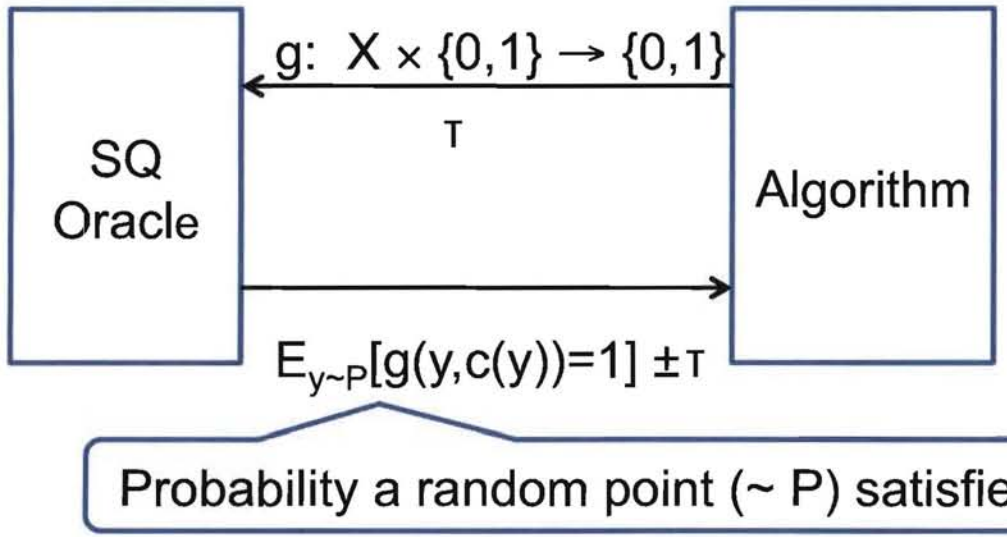

- ✓ $\tau > 1/\text{poly}(...)$
- ✓ g can be evaluated in poly time
- ✓ poly running time

Original goal: capture noise-resistant algorithms

- Most known learning problems fit this model
- Exception: Parity requires exponential # of SQ queries

Statistical Query (SQ) Model [Kearns93]

Same guarantees as PAC model, but algorithm no longer has access to individual labeled examples

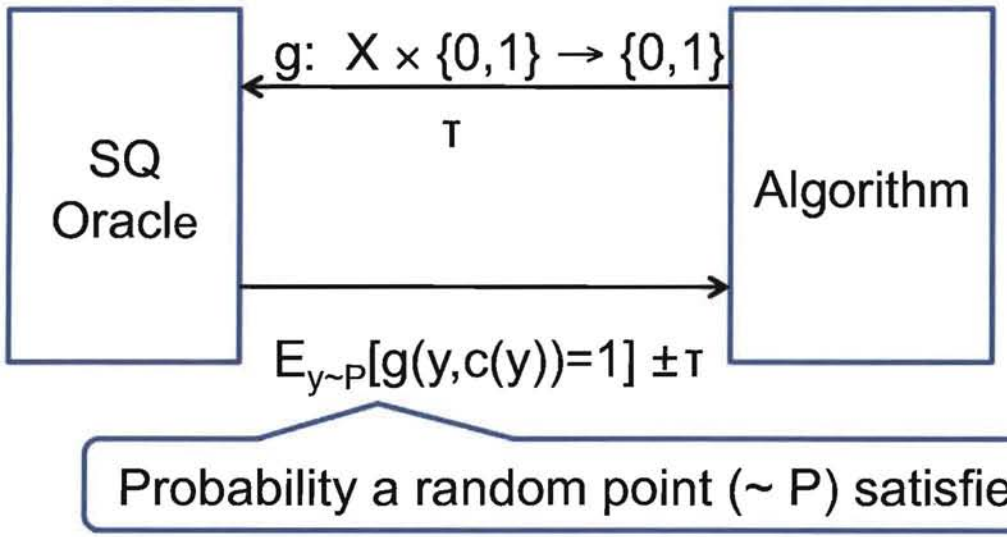

- ✓ $\tau > 1/\text{poly}(...)$
- ✓ g can be evaluated in poly time
- ✓ poly running time

I. Adaptive SQ: Queries can be adaptive
II. Non-adaptive SQ: Queries are fixed before the start

Our Results 3: Power of Interaction

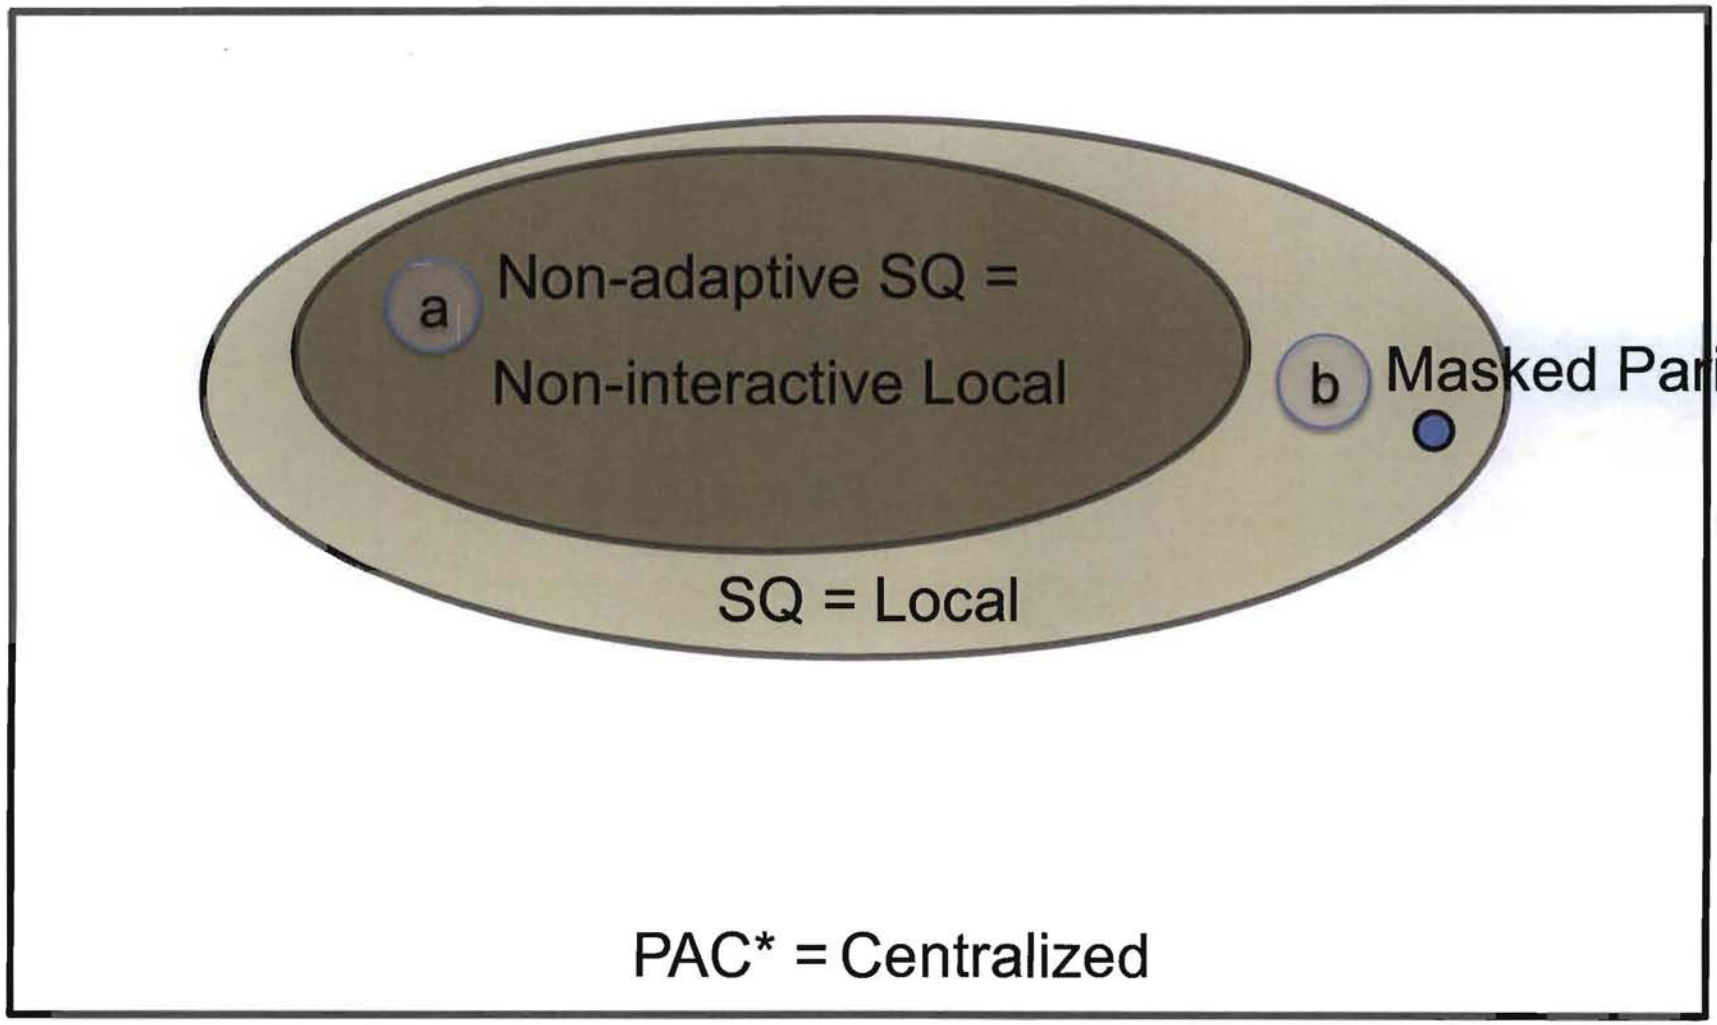

What Have We Learned

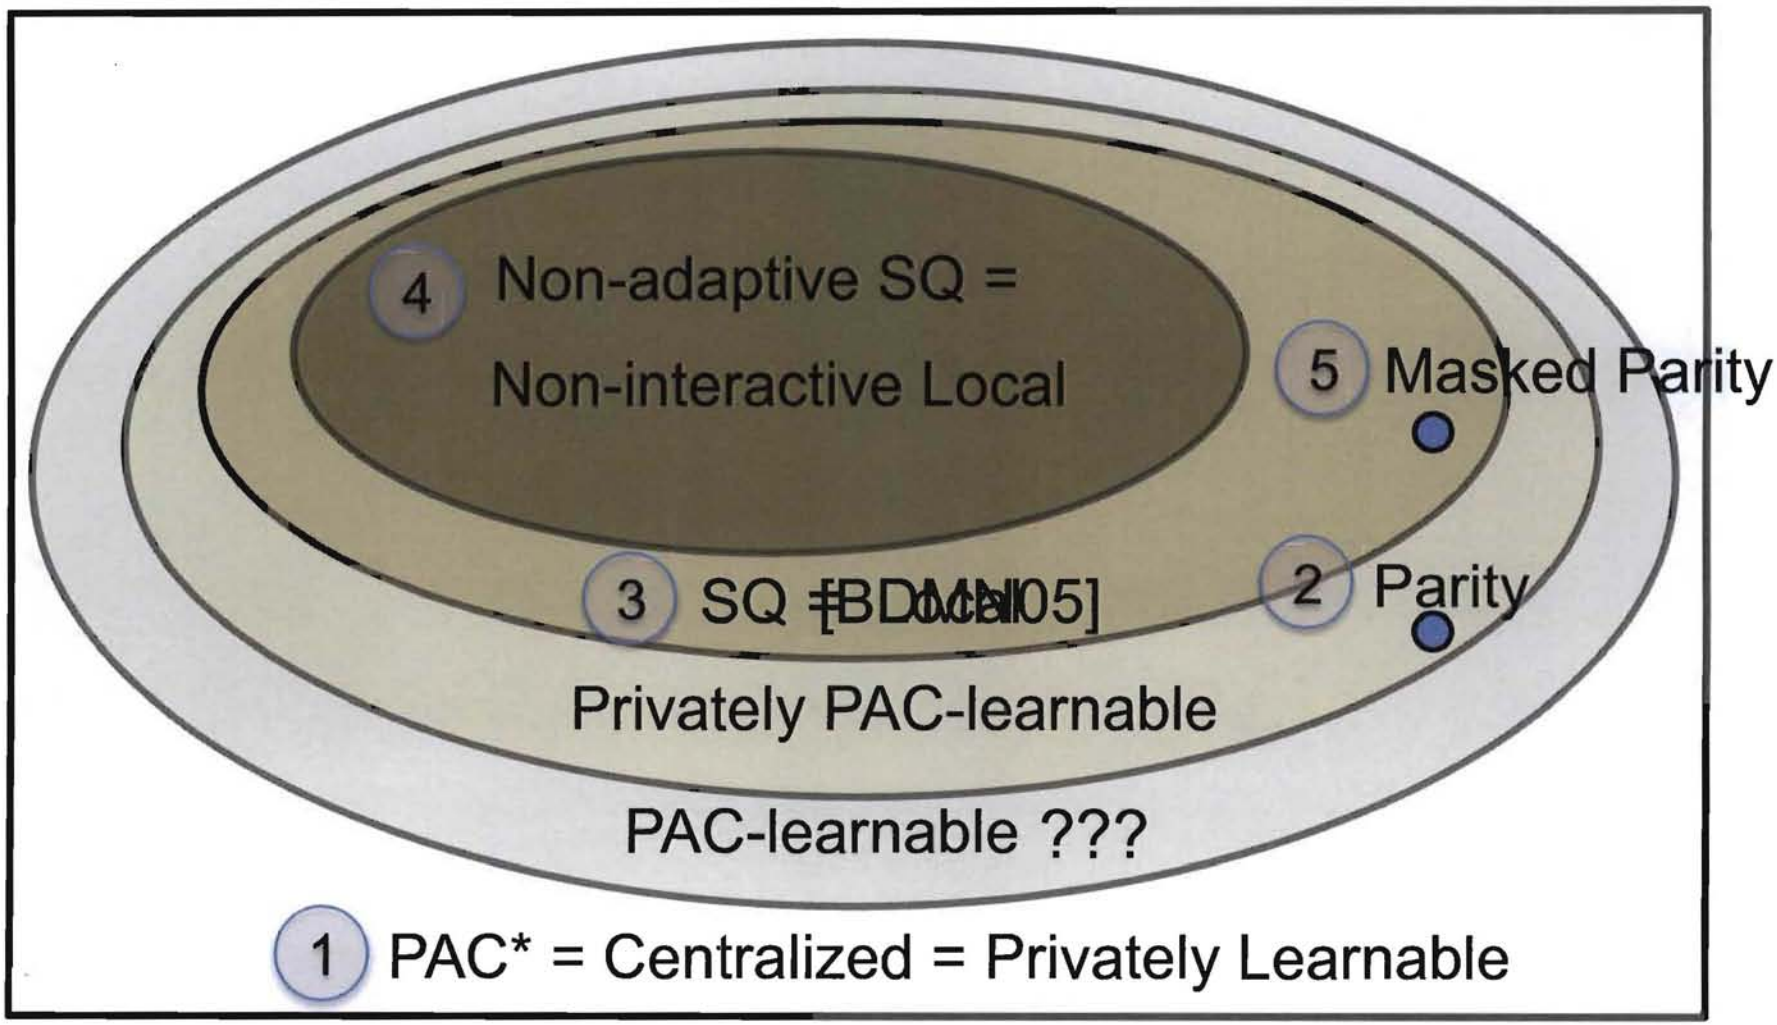

What We Want to Learn ???

- Separate efficient learning from efficient private learning
- Better private algorithms for SQ problems
- Other learning models

Thank You