2012 - 966-32D - Public Population Information in Differential Privacy - Christine Task AIP

CERIAS

The Center for Education and Research in Information Assurance and Security

PURDUE UNIVERSITY

Publicly Constrained Populations in Differential Privacy

Christine Task, Chris Clifton, Computer Science Department, Purdue University

Differentially private queries use noise to blur borders of neighboring possible worlds:

Q: "How many Purdue CS graduate students report seeking treatment for depression?"

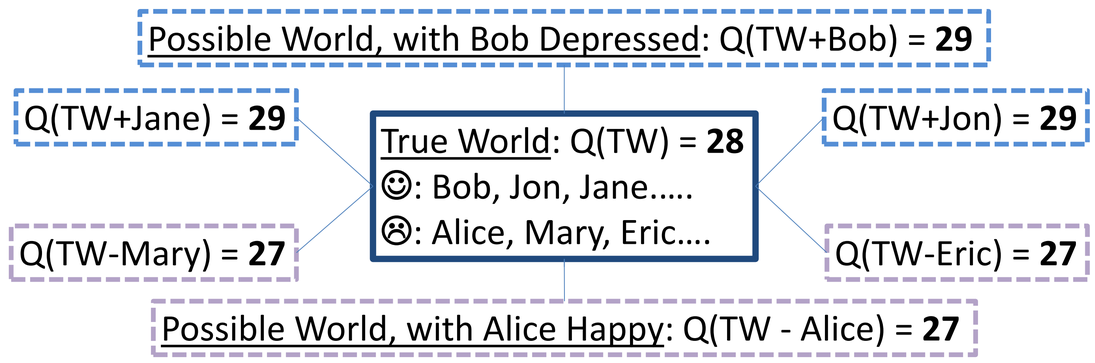

$Q_{private}$ (TW) = **27.2**
(Add random noise of standard deviation 1 to Q)
How are Alice, Bob? We can't tell: The result 27.2 could have come from any neighbor world.

Which pairs of worlds do we consider to be neighbors? Traditionally: Every pair of data-sets that differ in a single *arbitrary* individual data-point.

Query: "What is the teaching load of Purdue CS graduate students who report depression?" (results as histogram over # of labs taught)

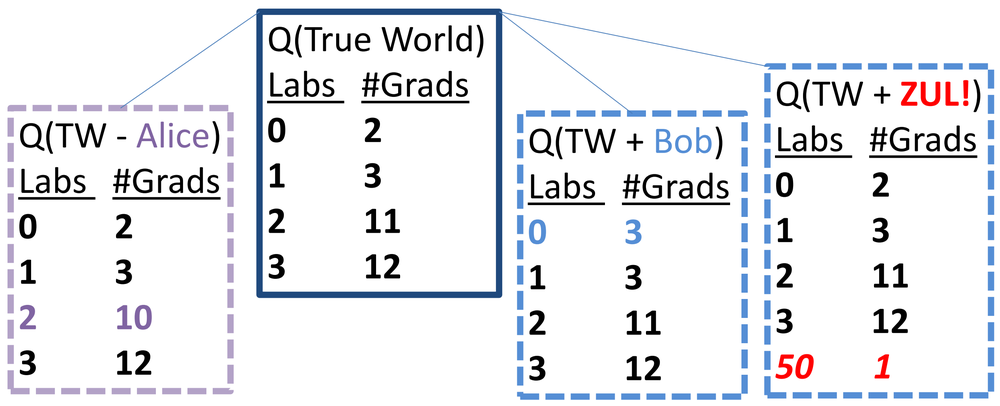

Q(TW - Alice)

Labs	#Grads
0	2
1	3
2	10
3	12

Q(True World)

Labs	#Grads
0	2
1	3
2	11
3	12

Q(TW + Bob)

Labs	#Grads
0	3
1	3
2	11
3	12

Q(TW + **ZUL!**)

Labs	#Grads
0	2
1	3
2	11
3	12
50	*1*

$Q_{private}$ (TW) =

Labs	#Grads
0	.9
1	5.1
2	10.7
3	13.3
…	
50	.8
51	.2
52	1.1
…	

What if it's public knowledge that some data-points simply *can't occur* in the data-set?

What does an attacker learn from the privatized query result *R*? For any neighbors *X, Y*:

$$\frac{Prob(True\ World\ (TW)=X\ |Q_{private}=R)}{Prob(True\ World\ (TW)=Y\ |Q_{private}=R)} = L_{x,y}\frac{Prob(TW=X)}{Prob(TW=y)},\ \text{where}\ L_{x,y} = \frac{Prob(Q_{pivate}=R\ |TW=X)}{Prob(Q_{pivate}=R\ |TW=Y)}$$

$L_{x,y}$ measures how much an attacker learns from *R* about the relative likelihood of *X* vs. *Y*: which world is more likely true. Differential Privacy requires that *L* be fairly close to 1, $L \le e^{\varepsilon}$ for small $\varepsilon > 0$. If $Q(X) - Q(Y)$ is large, then large noise is needed to ensure this.

But if *Y* has an arbitrary, impossible neighbor *Z*:

$$\frac{Prob(TW=Z\ |Q_{private}=R)}{Prob(TW=Y\ |Q_{private}=R)} = L_{z,y}\frac{0}{Prob(TW=y)} = 0$$

We can ignore $L_{z,y}$ when calibrating noise, because the attacker's belief in the relative likelihood of *X* vs. *Z* cannot change. She knows $Prob(Z) = 0$.

Many queries that appear to require too much noise to be feasibly privatizable in general, can be privatized with fairly little noise when public knowledge is considered.

Query	Generic Noise	Potential Context	Constrained Noise
How many triangles are there in graph *G*?	Unbounded	Given *G* is a subset of public school twitter follower graph that represents sexual relations.	Bounded by max individual triangle count of public graph
Average GPA in CS426?	Noise in denominator, orders of magnitude error	Given public knowledge that *n* people are registered for the class, no noise necessary in denominator.	Standard deviation of noise values is just $1/n$
How many in social network *S* have more than *k* friends?	Unbounded	Given that *S* limits members to at most *m* friends.	Standard deviation of noise is *m*, good when $\|S\| \gg m$

CERIAS

Discovery Park e-Enterprise Center