Engineering Privacy for Small Groups

Graham Cormode

g.cormode@warwick.ac.uk

Tejas Kulkarni (ATI/Warwick)

Divesh Srivastava (AT&T)

THE UNIVERSITY OF WARWICK

Many horror stories around data release...

Commonwealth of Massachusetts Group Insurance Commission
Your Benefits Connection

NETFLIX

NYC Taxi & Limousine Commission

AOL

We need to solve this data release problem...

Differential Privacy (Dwork et al 06)

A randomized algorithm K satisfies ε-differential privacy if:

Given two data sets that differ by one individual, D and D’, and any property S:

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

- Can achieve differential privacy for counts by adding a random noise value
- Uncertainty due to noise “hides” whether someone is present in the data

THE UNIVERSITY OF WARWICK

Achieving ε-Differential Privacy

(Global) Sensitivity of publishing:

$$s = \max_{x,x'} |F(x) - F(x')|, \; x, x' \text{ differ by 1 individual}$$

E.g., count individuals satisfying property P: one individual changing info affects answer by at most 1; hence $s = 1$

For every value that is output:

- Add Laplacian noise, $Lap(\varepsilon/s)$:

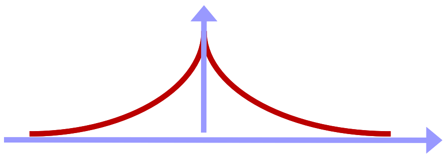

- Or Geometric noise for discrete case:

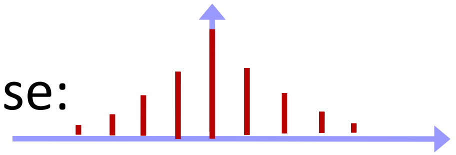

Simple rules for composition of differentially private outputs:

Given output O_1 that is ε_1 private and O_2 that is ε_2 private

- (Sequential composition) If inputs overlap, result is $\varepsilon_1 + \varepsilon_2$ private
- (Parallel composition) If inputs disjoint, result is $\max(\varepsilon_1, \varepsilon_2)$ private

Technical Highlights

- There are a number of building blocks for DP:
  - Geometric and Laplace mechanism for numeric functions
  - Exponential mechanism for sampling from arbitrary sets
    - Uses a user-supplied "quality function" for (input, output) pairs
- And "cement" to glue things together:
  - Parallel and sequential composition theorems
- With these blocks and cement, can build a lot
  - Many papers arrive from careful combination of these tools!
- Useful fact: any post-processing of DP output remains DP
  - (so long as you don't access the original data again)
  - Helps reason about privacy of data release processes

THE UNIVERSITY OF WARWICK

Limitations of Differential Privacy

- Differential privacy is NOT an algorithm but a property
  - Have to decide what algorithm to use and prove privacy properties
- Differential privacy does NOT guarantee utility
  - Naïve application of differential privacy may be useless
- The output of a differentially private process often does not have the same format as data input
- Basic model assumes that the data is held by a trusted aggregator

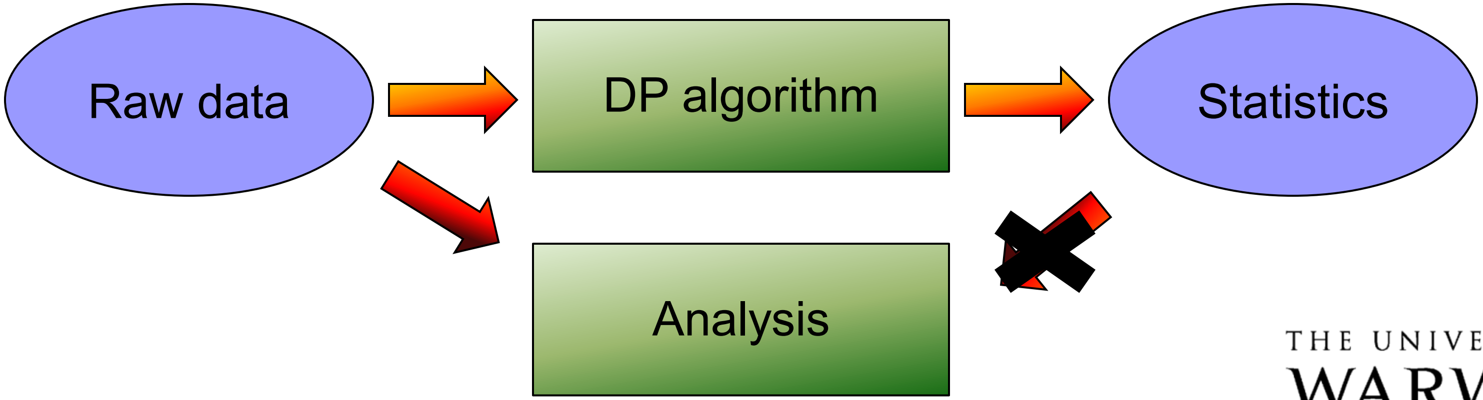

Local Differential Privacy

$$\begin{pmatrix} \mathrm{x} & x\alpha & x\alpha^2 & x\alpha^3 & \cdots & x\alpha^n \\ y\alpha & \mathrm{y} & y\alpha & y\alpha^2 & \cdots & y\alpha^{n-1} \\ y\alpha^2 & y\alpha & \mathrm{y} & y\alpha & \cdots & y\alpha^{n-2} \\ y\alpha^3 & y\alpha^2 & y\alpha & \mathrm{y} & \cdots & y\alpha^{n-3} \\ y\alpha^4 & y\alpha^3 & y\alpha^2 & y\alpha & \cdots & y\alpha^{n-4} \\ \vdots & \vdots & \vdots & \vdots & \ddots & \vdots \\ x\alpha^n & x\alpha^{n-1} & x\alpha^{n-2} & x\alpha^{n-3} & \cdots & \mathrm{x} \end{pmatrix}$$

- Data release under DP assumes a trusted third party aggregator
  - What if I don't want to trust a third party?
  - Use crypto?: fiddly secure multiparty computation protocols
- OR: run a DP algorithm with one participant for each user
  - Not as silly as it sounds: noise cancels over large groups
  - Implemented by Google and Apple (browsing/app statistics)
- Local Differential privacy state of the art in 2016: Randomized response (1965): five decade lead time!

- Lots of opportunity for new work:
  - Designing optimal mechanisms for local differential privacy
  - Adapt to apply beyond simple counts

Randomized Response and DP

- Developed as a technique for surveys with sensitive questions
  - "How will you vote in the election?"
  - Respondents may not respond honestly!
- Simple idea: tell respondents to lie (in a controlled way)
  - Randomized Response: Toss a coin with probability $p > \frac{1}{2}$
  - Answer truthfully if head, lie if tails
- Over a population of size n, expect $p\phi n + (1-p)(1-\phi)n$
  - Knowing p and n, solve for unknown parameter ϕ
- RR is DP: the ratio between the same output for different inputs is $p/(1-p)$
  - Larger p: more confidence (lower variance) but lower privacy
  - A local algorithm: no trusted aggregator

THE UNIVERSITY OF WARWICK

Small Group Privacy

- Many scenarios where there is a small group who trust each other with private data
  - A family who share a house
  - A team collaborating in an office
  - A group of friends in a social network

- They can gather their data together, and release through DP
  - Larger than the single entity model of local DP
  - But smaller than the general aggregation of data model
- We want to design *mechanisms* that have nice properties
  - A mechanism defines the output distribution, given the input

Mechanism Design

- We want to construct optimal mechanisms for data release
  - Target function: each user has a bit; release the sum of bits
  - Input range = output range = {0, 1, … n}
- Model a mechanism as a matrix of conditional probabilities $Pr[i|j]$
- DP introduces constraints on the matrix entries:

$$\alpha \, Pr[i|j] \leq Pr[i|j+1]$$

  - Neighbouring entries should differ by a factor of at most α
- We want to penalize outputs that are far from the truth:
  Define loss function $L_p = \sum_{i,j} w_j \, Pr[i|j] \, |i-j|^p * (n+1)/n$
  for weights (prior) w_j
  - We will focus on the core case of $p=0$, and uniform prior

THE UNIVERSITY OF WARWICK

Mechanism Properties

There are various properties we may want mechanisms to have:

- Row Honesty RH: $\forall$ i,j : Pr[i|i] $\geq$ Pr[i|j]
- Row Monotonicity RM: prob. decreases from Pr[i|i] along row
  - Row Monotonicity implies Row Honesty
- Column Honesty CH and Column Monotonicity CM, symmetrically
- Fairness F: $\forall$ i, j : Pr[i|i] = Pr[j|j]
  - Fairness and row honesty implies column honesty
- Weak honesty WH: Pr[i|i] $\geq$ 1/(n+1)
  - Achievable by the trivial uniform mechanism UM Pr[i|j] = 1/(n+1)
- Symmetry: $\forall$ i, j : Pr[i|j] = Pr[n-i|n-j]
  - Symmetry is achievable with no loss of objective function

Finding Optimal Mechanisms

- Goal: find optimal mechanisms for a given set of properties
- Can solve with optimization
  - Objective function is linear in the variables Pr[i|j]
  - Properties can all be specified as linear constraints on Pr[i|j]s
  - DP property is a linear constraint on Pr[i|j]s
- So can specify any desired set of combinations and solve an LP
- Patterns emerge… there are only a few distinct outcomes
  - Aim to understand the structure of optimal mechanisms
  - We seek explicit constructions
    - More efficient and amenable to analysis than solving LPs

Basic DP

- If we only seek DP, we always find a structured result
  - With symmetry and row monotonicity

$$\begin{pmatrix} \mathbf{x} & x\alpha & x\alpha^2 & x\alpha^3 & \cdots & x\alpha^n \\ y\alpha & \mathbf{y} & y\alpha & y\alpha^2 & \cdots & y\alpha^{n-1} \\ y\alpha^2 & y\alpha & \mathbf{y} & y\alpha & \cdots & y\alpha^{n-2} \\ y\alpha^3 & y\alpha^2 & y\alpha & \mathbf{y} & \cdots & y\alpha^{n-3} \\ y\alpha^4 & y\alpha^3 & y\alpha^2 & y\alpha & \cdots & y\alpha^{n-4} \\ \vdots & \vdots & \vdots & \vdots & \ddots & \vdots \\ x\alpha^n & x\alpha^{n-1} & x\alpha^{n-2} & x\alpha^{n-3} & \cdots & \mathbf{x} \end{pmatrix}$$

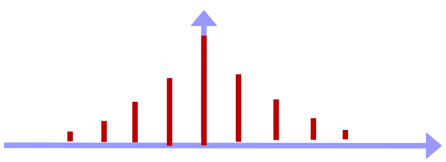

- Here $x = 1/(1+\alpha)$, $y=(1-\alpha)/(1+\alpha)$
- This is the truncated geometric mechanism GM [Ghosh et al. 09]:
  - Add symmetric geometric noise with parameter α to true answer
  - Truncate to range {0…n}
- Can prove this is the unique such optimal mechanism

Limitations of GM

- The Geometric Mechanism (GM) is not altogether satisfying
  - Tends to place a lot of weight on $\{0, n\}$ when α is large
- Misses most of the defined properties
  - Lacks Fairness ($Pr[i|i]=Pr[j|j]$)
  - Achieves Weak Honesty ($Pr[i|i]>Pr[i|j]$) only if $n > 2\alpha/(1-\alpha)$
  - Achieves Column Monotonicity only if $\alpha < ½$ (low privacy)
- But its L_0 score is the optimal value: $2\alpha / (1+\alpha)$
  - We seek more structured mechanisms that have similar score

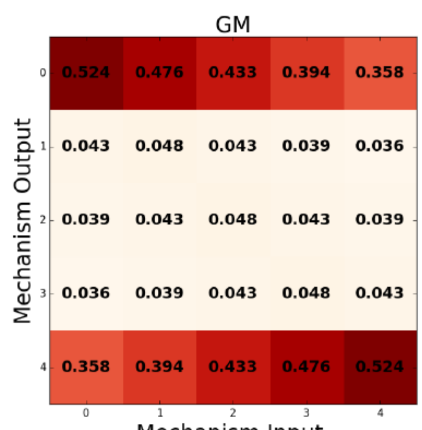

Example for $\alpha = 0.9$

THE UNIVERSITY OF WARWICK

Explicit Fair Mechanism EM

- We construct a new 'explicit fair mechanism' (uniform diagonal):

$$\begin{pmatrix} y & y\alpha & y\alpha^2 & y\alpha^3 & y\alpha^4 & y\alpha^4 & y\alpha^4 & y\alpha^4 \\ y\alpha & y & y\alpha & y\alpha^2 & y\alpha^3 & y\alpha^3 & y\alpha^3 & y\alpha^3 \\ y\alpha & y\alpha & y & y\alpha & y\alpha^2 & y\alpha^3 & y\alpha^3 & y\alpha^3 \\ y\alpha^2 & y\alpha^2 & y\alpha & y & y\alpha & y\alpha^2 & y\alpha^2 & y\alpha^2 \\ y\alpha^2 & y\alpha^2 & y\alpha^2 & y\alpha & y & y\alpha & y\alpha^2 & y\alpha^2 \\ y\alpha^3 & y\alpha^3 & y\alpha^3 & y\alpha^2 & y\alpha & y & y\alpha & y\alpha \\ y\alpha^3 & y\alpha^3 & y\alpha^3 & y\alpha^3 & y\alpha^2 & y\alpha & y & y\alpha \\ y\alpha^4 & y\alpha^4 & y\alpha^4 & y\alpha^4 & y\alpha^3 & y\alpha^2 & y\alpha & y \end{pmatrix}$$

  - Each column is a permutation of the same set of values
  - Additionally has column and row monotonicity, symmetry
  - This is an optimal fair mechanism:
    - Entries in middle column are all as small as DP will allow
    - Hence y cannot be bigger
    - Cost slightly higher than Geometric Mechanism

Summary of mechanisms

- Based on relations between properties, we can conclude:

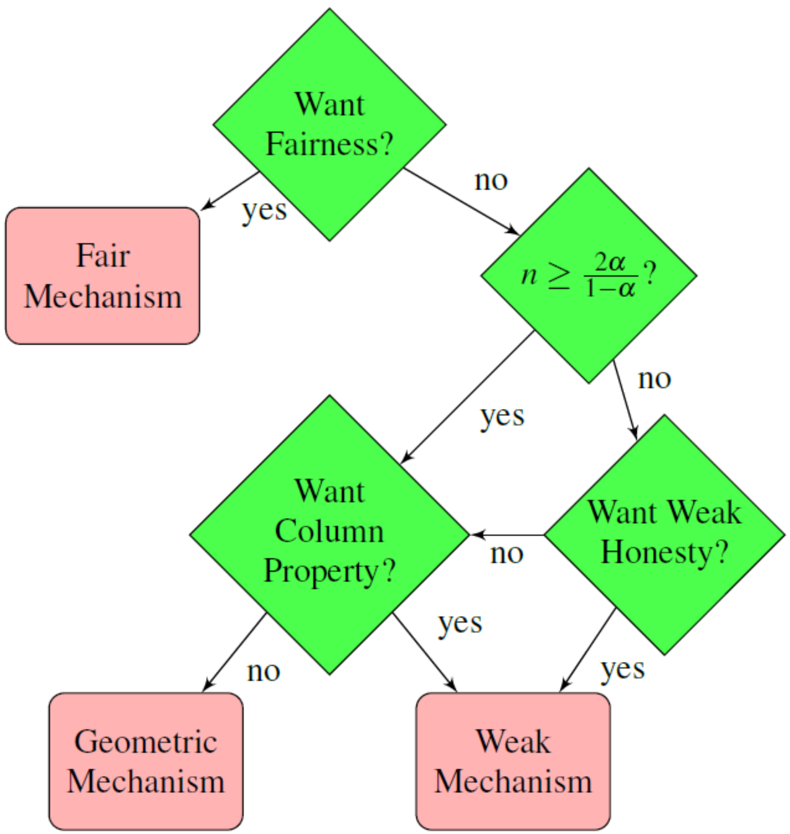

- Fair Mechanism (EM) and Geometric Mechanism (GM) have explicit forms
- Weak Mechanism (WM) found by solving LP with weak honesty constraint

Property	GM	UM	EM	WM
Symmetry (S)	Y	Y	Y	Y
Row Monotone (RM)	Y	Y	Y	Y
Column Monotone (CM)	—	Y	Y	Y
Fairness (F)	N	Y	Y	N
Weak Honesty (WH)	—	Y	Y	Y
$\mathbb{L}_0$	$\frac{2\alpha}{1+\alpha}$	1	$\approx \frac{2\alpha}{1+\alpha} \cdot \frac{n+1}{n}$	$\geq \frac{2\alpha}{1+\alpha}$

Comparing Mechanisms

- Heatmaps comparing mechanisms for $\alpha = 0.9$, n=4

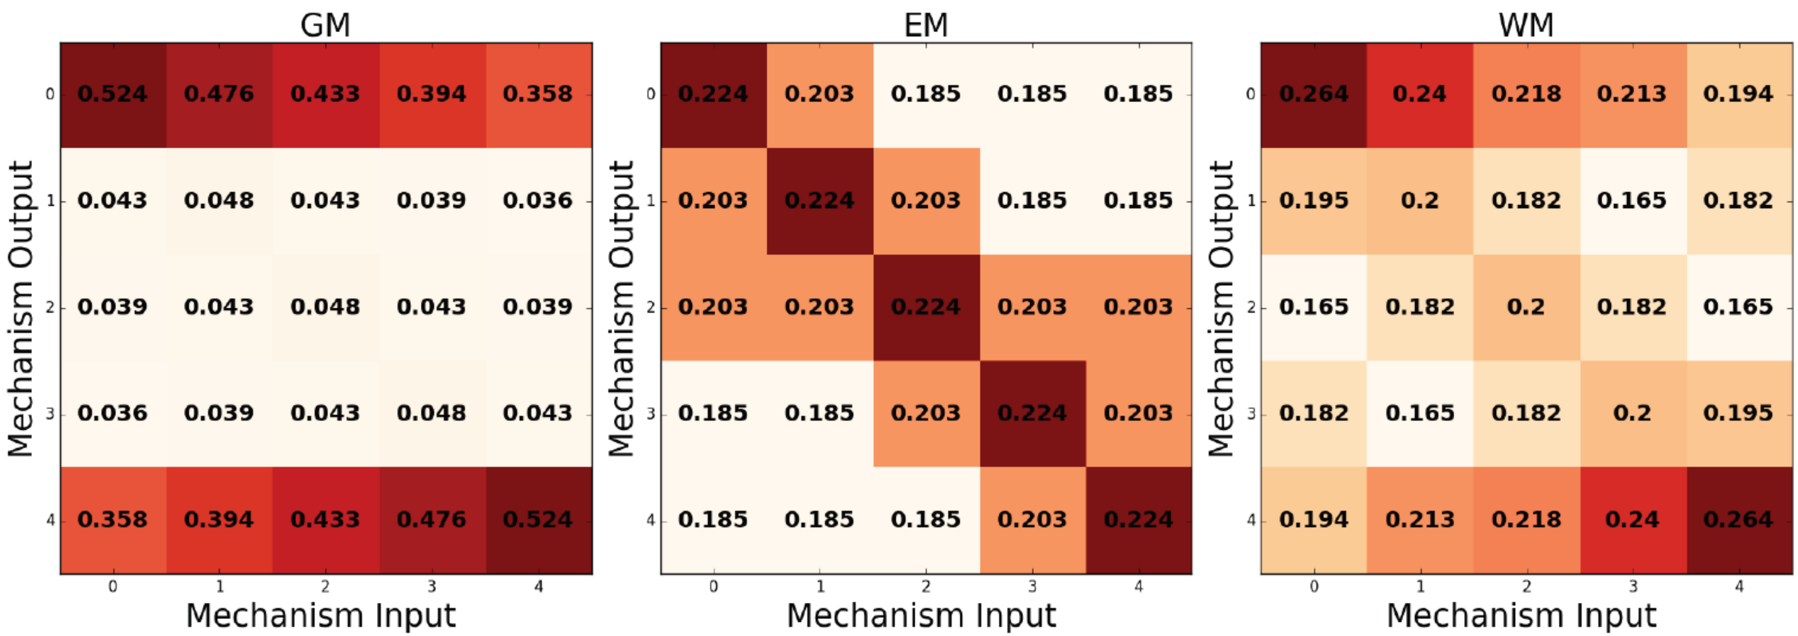

THE UNIVERSITY OF WARWICK

L_0 score behaviour

- L_0 score varies as a function of n and α
  - WM converges on GM for $n \geq 2\alpha / (1-\alpha)$

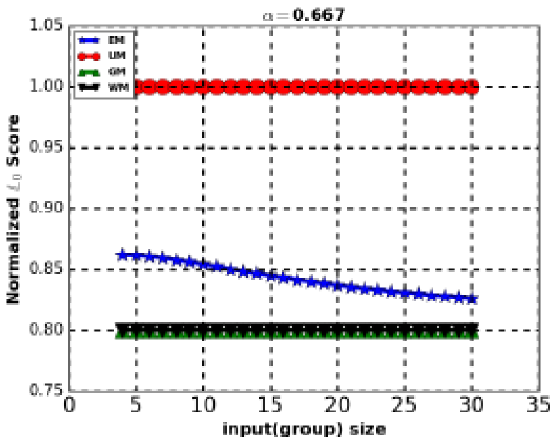

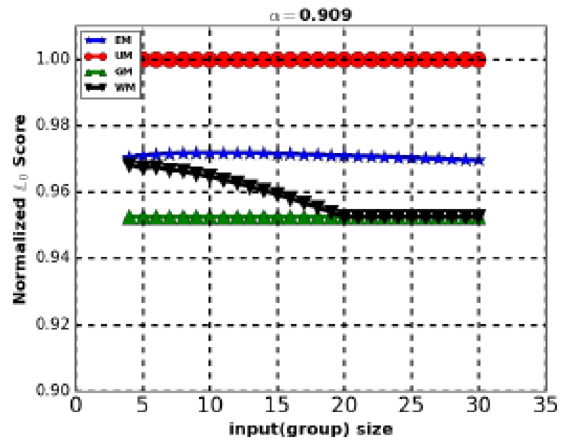

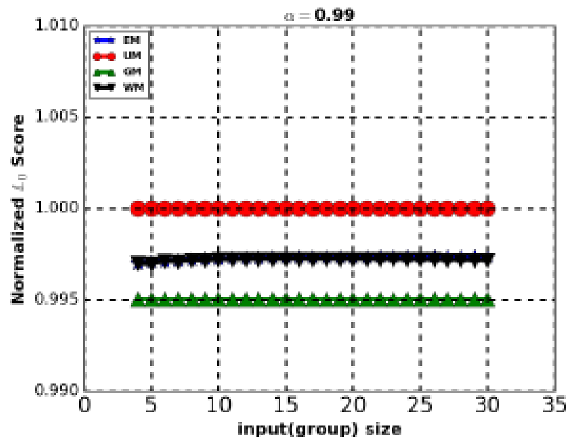

THE UNIVERSITY OF WARWICK

Performance on real data

- Using UCI Adult data set of demographic data
  - Construct small groups in the data, target different binary attributes
  - Compute Root-Mean-Squared Error of per-group outputs
  - EM and WM generally preferable for wide range of α values

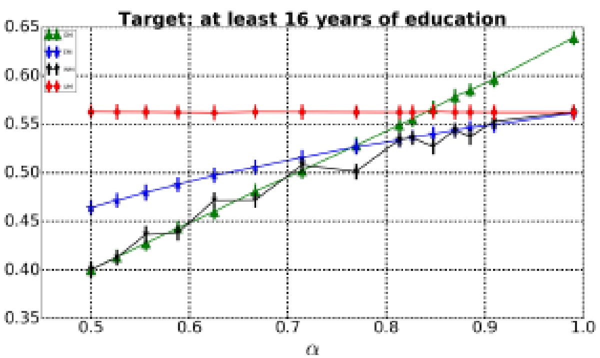

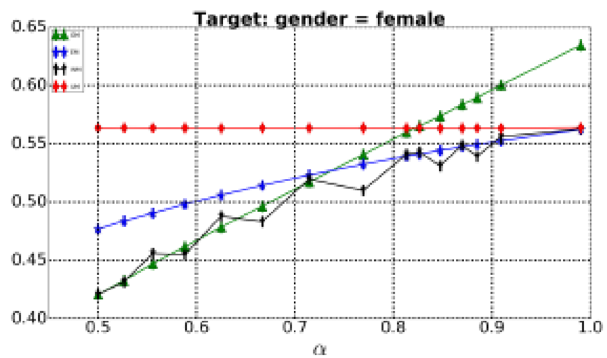

THE UNIVERSITY OF WARWICK

Summary

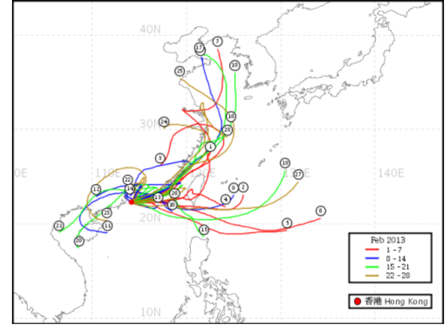

- Carefully crafted mechanisms for data release perform well on small groups
- Many more natural questions for small groups and local DP
- Lots of technical work left to do:
  - Structured data: other statistics, graphs, movement patterns
  - Unstructured data: text, images, video?
  - Develop standards for (certain kinds of) data release

Joint work with Divesh Srivastava (AT&T), Tejas Kulkarni (Warwick)
Supported by AT&T, Royal Society, European Commission

THE UNIVERSITY OF WARWICK