Constrained Private Mechanisms for Count Data

Graham Cormode

g.cormode@warwick.ac.uk

Tejas Kulkarni (ATI/Warwick)

Divesh Srivastava (AT&T)

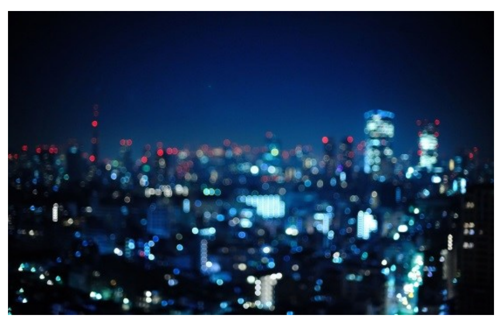

THE UNIVERSITY OF WARWICK

Private Data Release

- Many problems require collection of aggregate data
  - Simple count queries for statistics
  - Frequency parameters of analytic models

- The model of Differential Privacy (DP) gives a rigorous statistical definition
  - Requires each output to have a similar probability as inputs vary
- Our aim is to design *mechanisms* that have nice properties
  - A mechanism defines the output distribution, given the input
  - We seek accurate, usable outputs, from small groups

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 = $\{0, 1, \ldots n\}$
- Model a mechanism as a matrix of conditional probabilities $Pr[i|j]$
- DP introduces constraints on the matrix entries:

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

  - Neighboring entries should differ by a factor of at most $0 < \alpha < 1$
- We want to penalize outputs that are far from the truth:
  Define loss function $L_p = \sum_{i,j} w_j \, Pr[i|j] \, |i - j|^p$ for weights (prior) w_j
  - We will focus on the core case of $p=0$, and uniform prior
  - L_0 loss function is then just the sum of weights off-diagonal
  - Equivalently, maximize trace of the probability matrix

Unconstrained Mechanism: GM

- Optimizing for L_0 loss function yields a highly structured result:

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

- Here $x = 1/(1+\alpha)$, $y=(1-\alpha)/(1+\alpha)$, $L_0=2\alpha/(1+\alpha)$

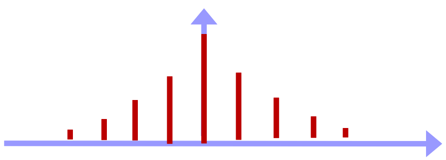

- This is the truncated geometric mechanism **GM** [Ghosh et al. 09]:
  - Add symmetric geometric noise with parameter α to true answer
  - Truncate to range $\{0\ldots n\}$
- We prove this is the unique such optimal mechanism for L_0
  - But it has some issues!

Limitations of GM

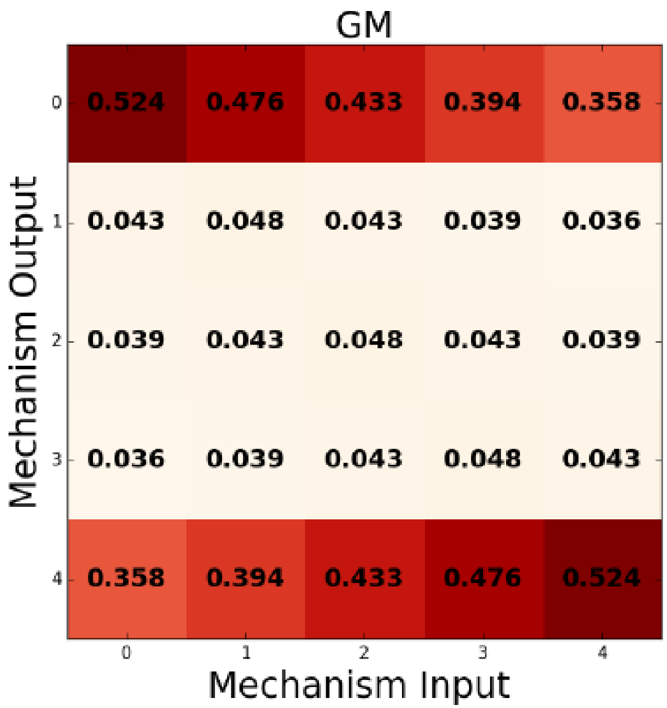

Example for
$\alpha = 0.9$

- GM tends to place a lot of weight on $\{0, n\}$ when α is large
  - But GM's L_0 score is the optimal value: $2\alpha / (1+\alpha)$
  - The issue is even worse if optimizing for L_1 or L_2 objective functions!
  - We seek more structured mechanisms that have similar score

Mechanism Properties

We give 7 constraints to impose more structure on mechanisms:

- Row Honesty RH: $\forall$ i,j : Pr[i|i] $\geq$ Pr[i|j] (true value is most likely)
- 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] (same probability of truthfulness)
  - Fairness and row honesty implies column honesty
- Weak honesty WH: Pr[i|i] $\geq$ 1/(n+1) (at least uniformly truthful)
  - 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 techniques
  - 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
  - Always feasible: just uniform guessing (UM) meets all constraints
- Patterns emerge: of 127 possibilities, only few distinct outcomes
  - Aim to understand the structure of optimal mechanisms
  - We seek explicit constructions
    - More efficient and amenable to analysis than solving LPs

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
- Has all our properties: column & row monotonicity, symmetry
- This is (one) 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:

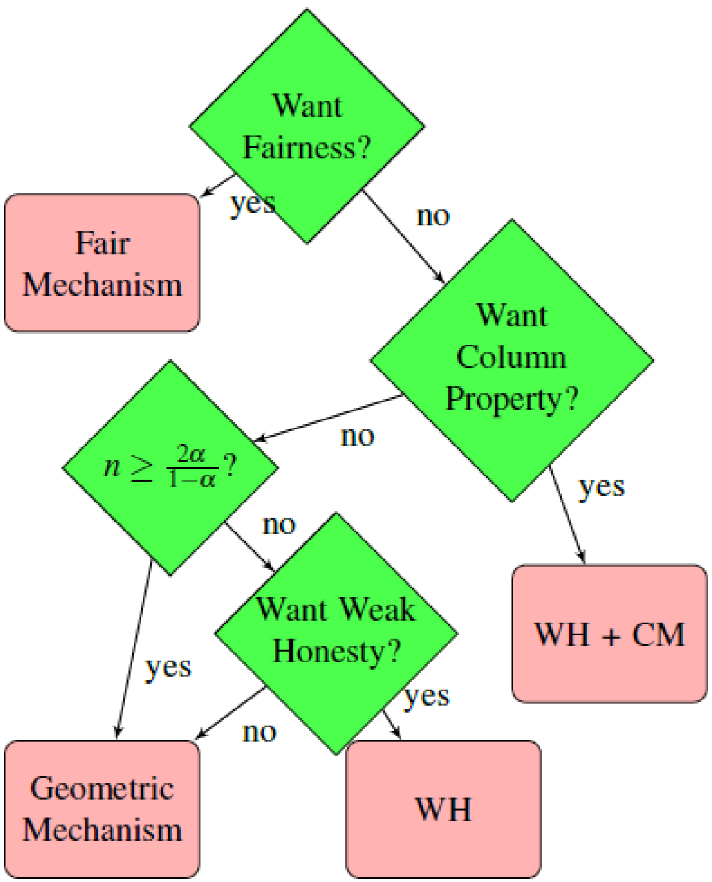

- Fair Mechanism (EM) and Geometric Mechanism (GM) have explicit forms
- Two Weak Mechanism variants (WM) found by solving LPs

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	$\approx \frac{2\alpha}{1+\alpha} \cdot \frac{n+1}{n}$	$\geq \frac{2\alpha}{1+\alpha}$

THE UNIVERSITY OF WARWICK

Comparing Mechanisms

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

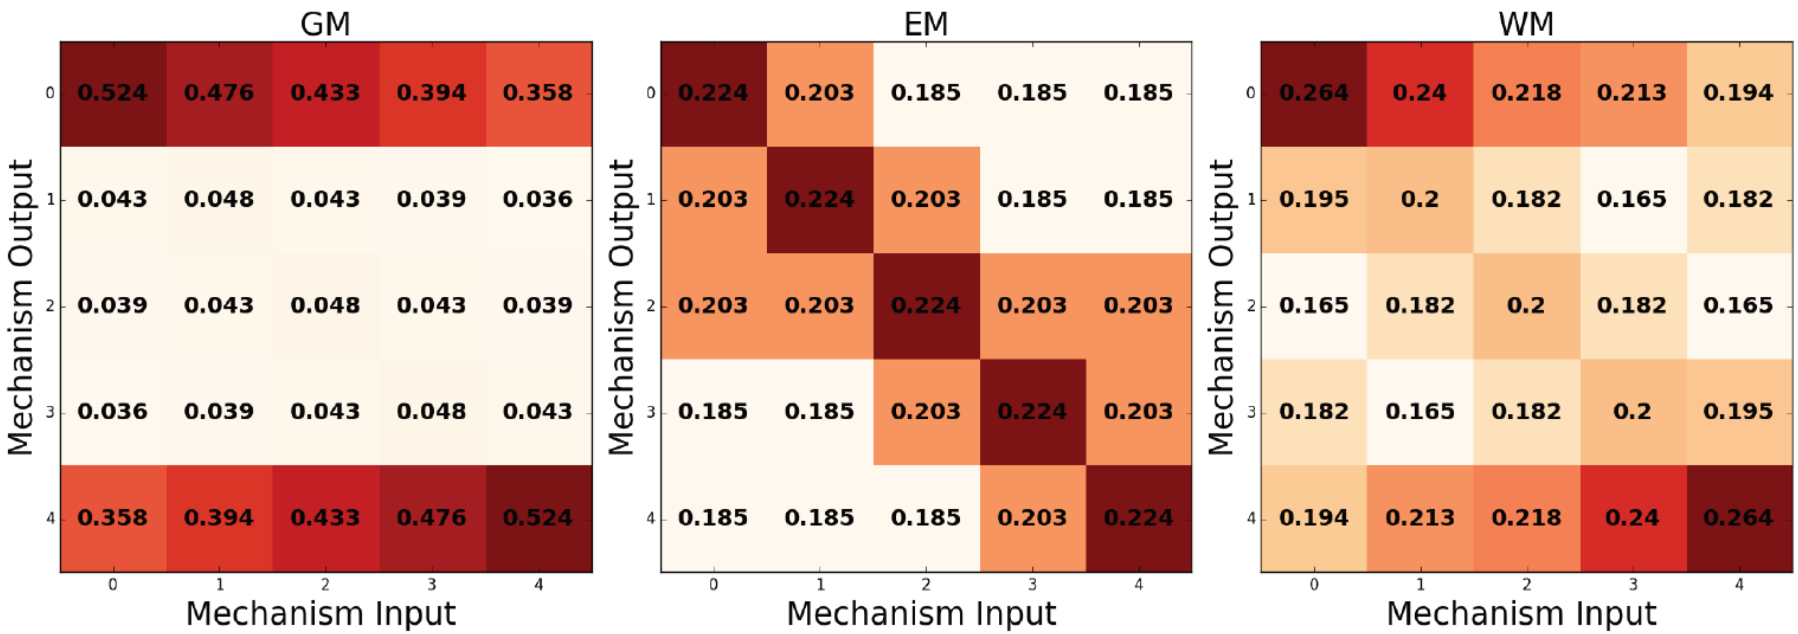

- Heatmaps look very different but their L_0 scores are close:

	GM	EM	WM
L_0 score	0.764	0.776	0.774

L_0 score behaviour

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

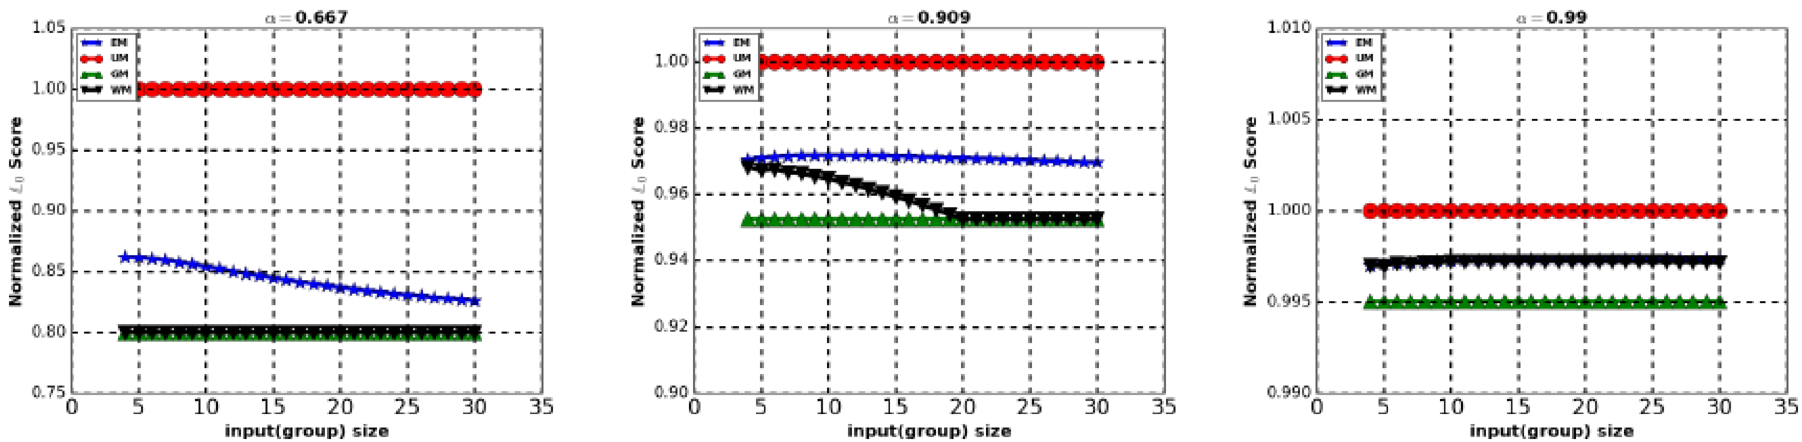

THE UNIVERSITY OF WARWICK

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

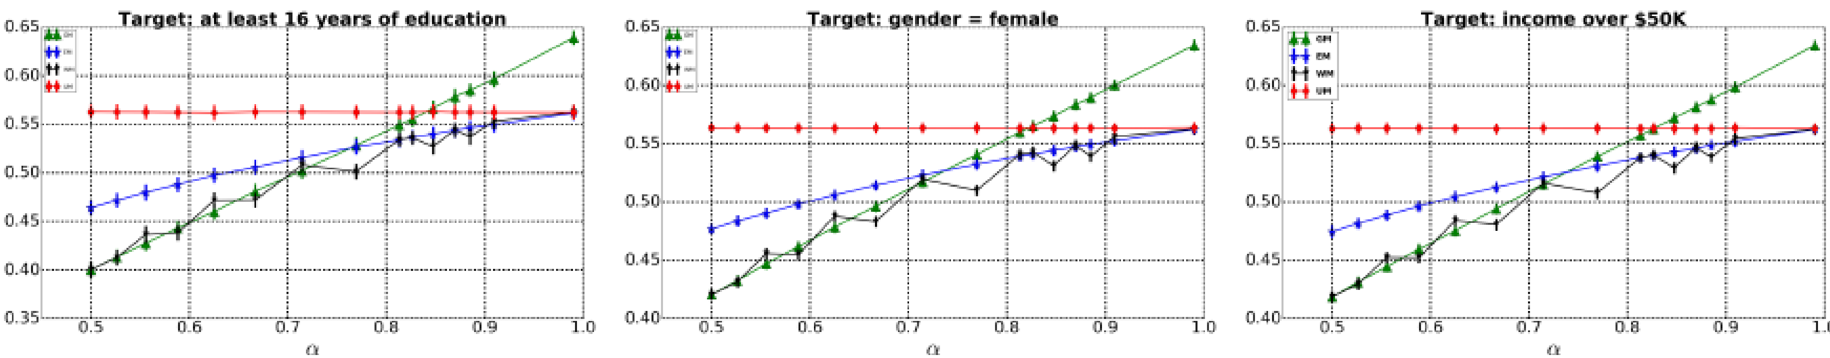

THE UNIVERSITY OF WARWICK

Summary

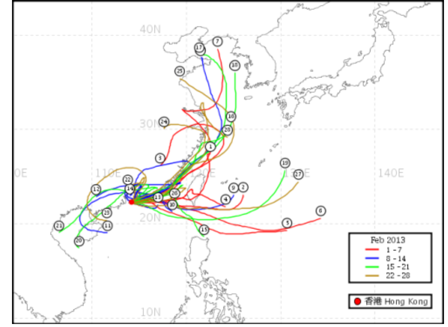

- Carefully crafted mechanisms for data release can fix anomalies/unexpected behavior for small groups
- Many more natural questions for small groups
  - Interpret constraints as regularization
  - Find closed form solutions for other objective functions (L_1, L_2)
- More general data release problems:
  - Structured data: other statistics, graphs, movement patterns
  - Unstructured data: text, images, video?

Supported by AT&T, Warwick, Alan Turing Institute, European Commission