What Does The Crowd Say About You? Evaluating Aggregation-Based Location Privacy

PETS 2017

Apostolos Pyrgelis [1], Carmela Troncoso [2] and Emiliano De Cristofaro [1]

[1]University College London, [2]IMDEA Software Institute

July 19, 2017
Minneapolis, USA

Table of Contents

Motivation

Motivation

- Aggregate locations are useful for performing mobility analytics, e.g., predicting traffic volumes, detecting anomalies

Motivation

- Aggregate locations are useful for performing mobility analytics, e.g., predicting traffic volumes, detecting anomalies
- Aggregation is often considered a privacy-friendly approach, especially when done in a privacy-preserving way

Motivation

- Aggregate locations are useful for performing mobility analytics, e.g., predicting traffic volumes, detecting anomalies
- Aggregation is often considered a privacy-friendly approach, especially when done in a privacy-preserving way
- **Open Problem:** No sound methodology to reason about privacy leakage for individuals from the aggregates

Motivation

- Aggregate locations are useful for performing mobility analytics, e.g., predicting traffic volumes, detecting anomalies
- Aggregation is often considered a privacy-friendly approach, especially when done in a privacy-preserving way
- **Open Problem:** No sound methodology to reason about privacy leakage for individuals from the aggregates
- ...no way to evaluate potential defense mechanisms

In this work...

In this work...

- We introduce a framework to address this gap

In this work...

- We introduce a framework to address this gap
- We use it to measure users' **privacy loss** from aggregate location time-series

In this work...

- We introduce a framework to address this gap
- We use it to measure users' **privacy loss** from aggregate location time-series
- We evaluate two real-world mobility datasets (TFL, SFC)

In this work...

- We introduce a framework to address this gap
- We use it to measure users' **privacy loss** from aggregate location time-series
- We evaluate two real-world mobility datasets (TFL, SFC)
- We study the privacy protection offered by defense mechanisms based on DP (output and input perturbation)

In this work...

- We introduce a framework to address this gap
- We use it to measure users' **privacy loss** from aggregate location time-series
- We evaluate two real-world mobility datasets (TFL, SFC)
- We study the privacy protection offered by defense mechanisms based on DP (output and input perturbation)
- Raw aggregates do leak information about users' locations and mobility profiles

In this work...

- We introduce a framework to address this gap
- We use it to measure users' **privacy loss** from aggregate location time-series
- We evaluate two real-world mobility datasets (TFL, SFC)
- We study the privacy protection offered by defense mechanisms based on DP (output and input perturbation)
- Raw aggregates do leak information about users' locations and mobility profiles
- DP provides strong privacy protection when the utility of the aggregates is poor

Table of Contents

Adversarial Prior Knowledge

Adversarial Prior Knowledge

- Prior knowledge might come from social networks, data leaks, location traces released by providers or personal knowledge, e.g. home / work pair

Adversarial Prior Knowledge

- Prior knowledge might come from social networks, data leaks, location traces released by providers or personal knowledge, e.g. home / work pair
- In this work, we explore several possible approaches

Adversarial Prior Knowledge

- Prior knowledge might come from social networks, data leaks, location traces released by providers or personal knowledge, e.g. home / work pair
- In this work, we explore several possible approaches
- **Probabilistic:**
  - Frequency of locations (over time)
  - Location Seasonality (day / week)

Adversarial Prior Knowledge

- Prior knowledge might come from social networks, data leaks, location traces released by providers or personal knowledge, e.g. home / work pair
- In this work, we explore several possible approaches
- **Probabilistic:**
  - Frequency of locations (over time)
  - Location Seasonality (day / week)
- **Assignment:**
  - Most popular locations
  - All prior locations
  - Last Season (hour / day / week)

Inference Strategies

Inference Strategies

- **Bayesian Update:**
  - Posterior probability of a user being in a location at a certain time, given the prior and the aggregates

Inference Strategies

- **Bayesian Update:**
  - Posterior probability of a user being in a location at a certain time, given the prior and the aggregates
- **MAX-ROI:**
  - Greedy strategy that assigns the most probable users to each location, until the aggregates are consumed

Inference Strategies

- **Bayesian Update:**
  - Posterior probability of a user being in a location at a certain time, given the prior and the aggregates
- **MAX-ROI:**
  - Greedy strategy that assigns the most probable users to each location, until the aggregates are consumed
- **MAX-USER:**
  - Greedy strategy that assigns each user to her most likely locations, until the aggregates are consumed

Timeline

Timeline

- **Observation Period:** used to build prior knowledge for each user

Timeline

- **Observation Period:** used to build prior knowledge for each user
- **Inference period:** launch the attacks

Adversarial Goals

Adversarial Goals

User Profiling:

- Infer the probability of a user being in a location at a certain time
- Adversarial Error: JS-divergence from the ground truth

Adversarial Goals

User Profiling:

- Infer the probability of a user being in a location at a certain time
- Adversarial Error: JS-divergence from the ground truth

User Localization:

- Predict where the user will be at a certain time
- Adversarial Error: 1 - F1

Privacy Loss

Privacy Loss

Privacy Loss (PL):

- normalized reduction in adversarial error with vs without the aggregates

Table of Contents

Datasets

Datasets

Tranport For London (TFL):

- 60M trips - 4M unique oyster cards - 582 stations (regions of interest - ROIs)
- Monday, March 1 - Sunday, March 28, 2010
- Sample the top 10K oyster ids per total # of trips, being active for 115 ± 21 out of the 672 timeslots and reporting 171 ± 26 ROIs in total (sparse, regular)

Datasets

Tranport For London (TFL):

- 60M trips - 4M unique oyster cards - 582 stations (regions of interest - ROIs)
- Monday, March 1 - Sunday, March 28, 2010
- Sample the top 10K oyster ids per total # of trips, being active for 115 ± 21 out of the 672 timeslots and reporting 171 ± 26 ROIs in total (sparse, regular)

San Francisco Cabs (SFC):

- 11M GPS coordinates - 534 cabs in SF - May 19 to June 8, 2008
- Grid $10 \times 10 = 100$ ROIs of 0.5×0.37 mi^2
- Taxis are active for 340 ± 94 out of the 504 timeslots and report $3,663 \pm 1,116$ ROIs in total (dense, irregular)

Datasets

Tranport For London (TFL):

- 60M trips - 4M unique oyster cards - 582 stations (regions of interest - ROIs)
- Monday, March 1 - Sunday, March 28, 2010
- Sample the top 10K oyster ids per total # of trips, being active for 115 ± 21 out of the 672 timeslots and reporting 171 ± 26 ROIs in total (sparse, regular)

San Francisco Cabs (SFC):

- 11M GPS coordinates - 534 cabs in SF - May 19 to June 8, 2008
- Grid $10 \times 10 = 100$ ROIs of 0.5×0.37 mi^2
- Taxis are active for 340 ± 94 out of the 504 timeslots and report $3,663 \pm 1,116$ ROIs in total (dense, irregular)

Prior Knowledge: Split data according to time

Datasets

Tranport For London (TFL):

- 60M trips - 4M unique oyster cards - 582 stations (regions of interest - ROIs)
- Monday, March 1 - Sunday, March 28, 2010
- Sample the top 10K oyster ids per total # of trips, being active for 115 ± 21 out of the 672 timeslots and reporting 171 ± 26 ROIs in total (sparse, regular)

San Francisco Cabs (SFC):

- 11M GPS coordinates - 534 cabs in SF - May 19 to June 8, 2008
- Grid $10 \times 10 = 100$ ROIs of 0.5×0.37 mi^2
- Taxis are active for 340 ± 94 out of the 504 timeslots and report $3,663 \pm 1,116$ ROIs in total (dense, irregular)

Prior Knowledge: Split data according to time
Inference Period: Last week of each dataset

User Profiling

User Profiling

Prior Knowledge: Location frequency, over the observation period

Prior Knowledge: Location frequency, over the observation period

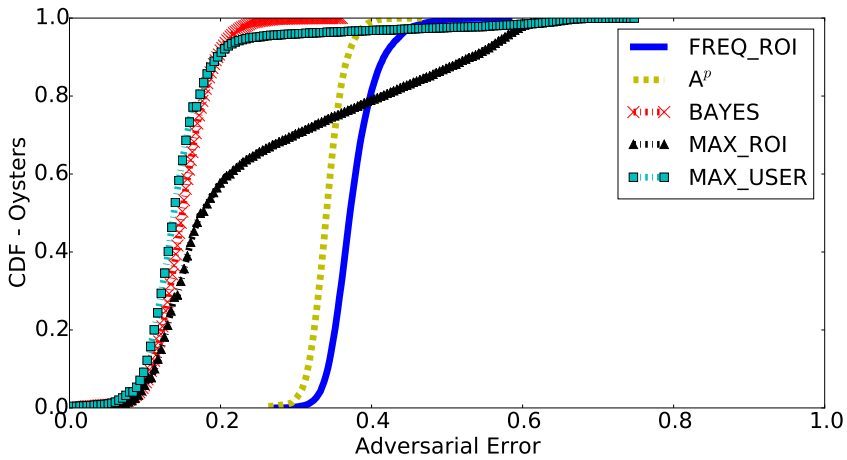

TFL

Prior Knowledge: Location frequency, over the observation period

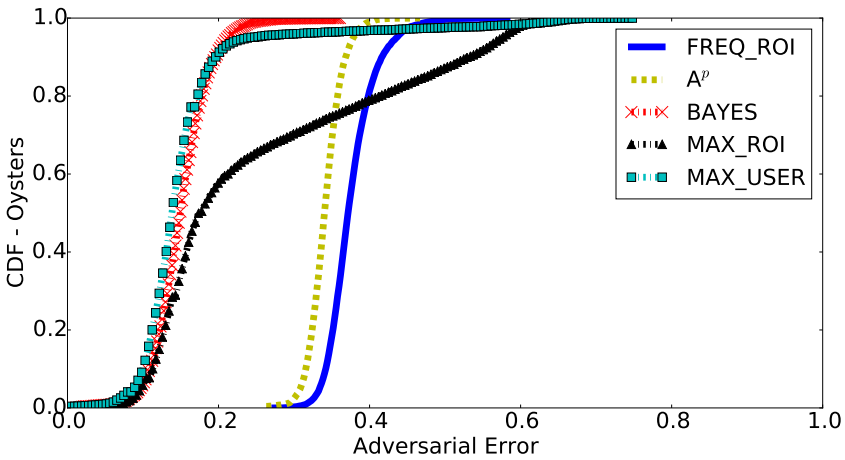

TFL

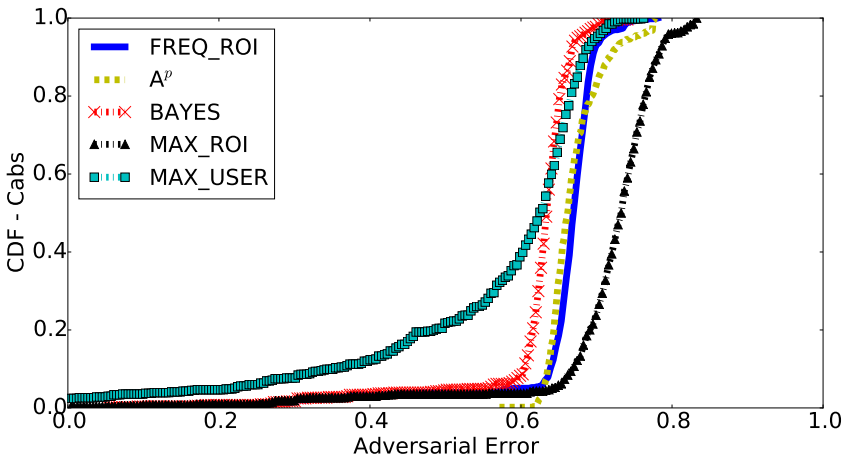

SFC

User Profiling

Prior Knowledge: Location frequency, over the observation period

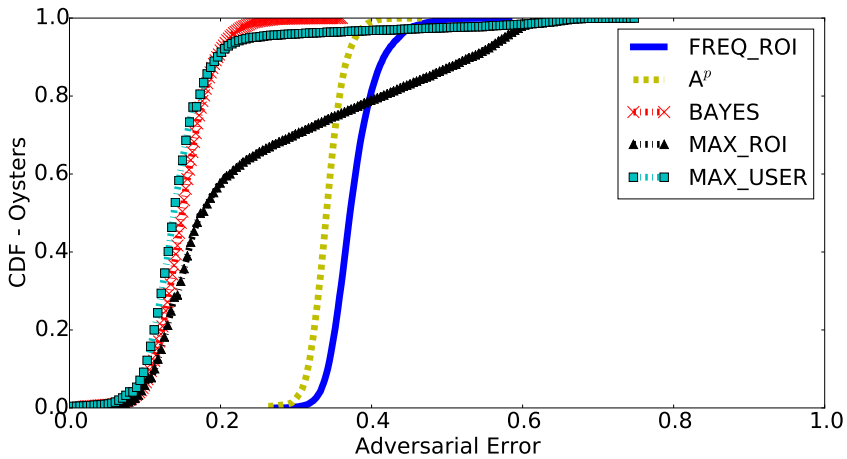

TFL

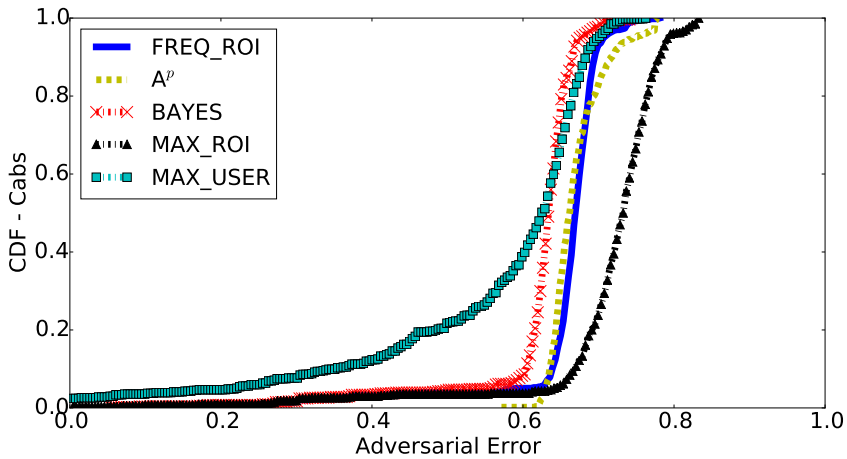

SFC

TFL vs. SFC: Inferring mobility profiles of commuters easier than cabs

Implications of Regular Mobility Patterns

Implications of Regular Mobility Patterns

Prior Knowledge: Location frequency, for time instances of any day

Implications of Regular Mobility Patterns

Prior Knowledge: Location frequency, for time instances of any day

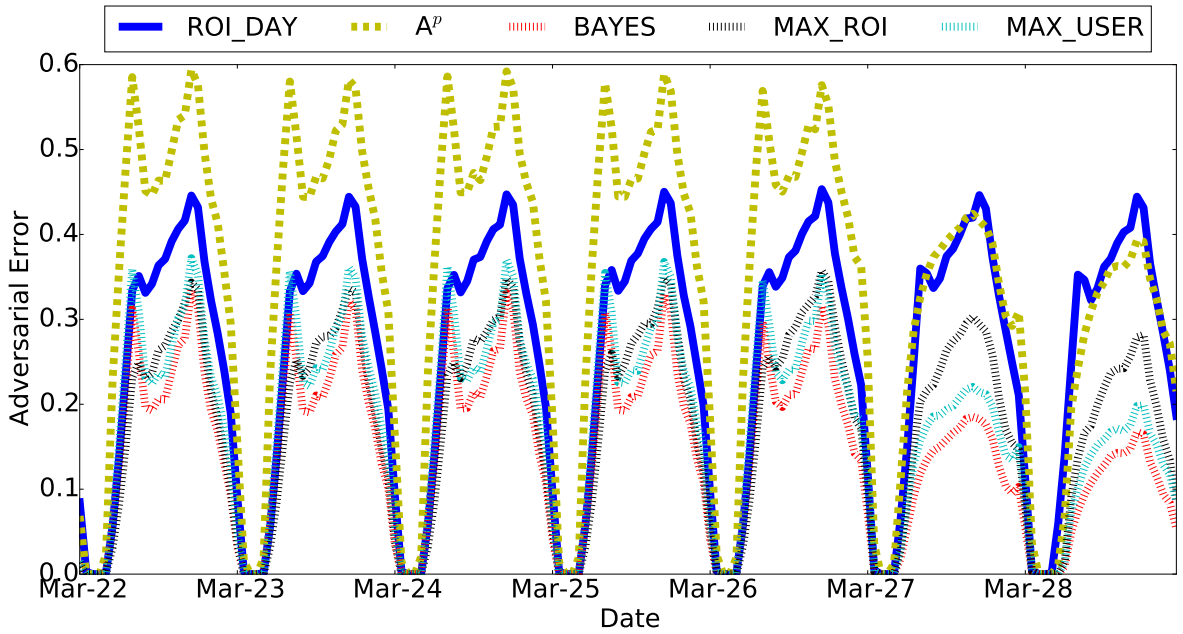

TFL

User Localization

User Localization

Prior Knowledge: Location frequency, for time instances of a week

User Localization

Prior Knowledge: Location frequency, for time instances of a week

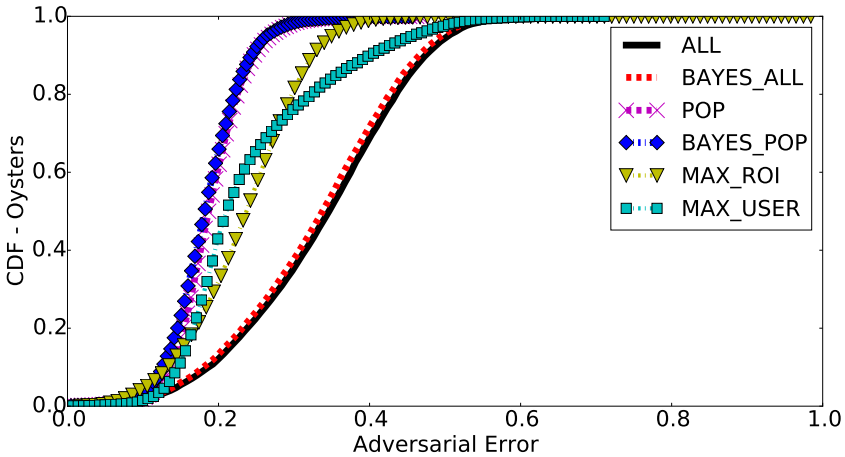

TFL

User Localization

Prior Knowledge: Location frequency, for time instances of a week

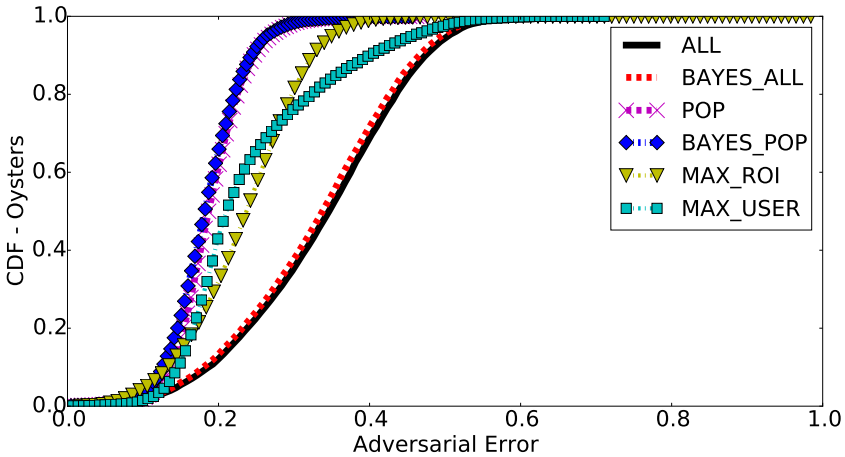

TFL

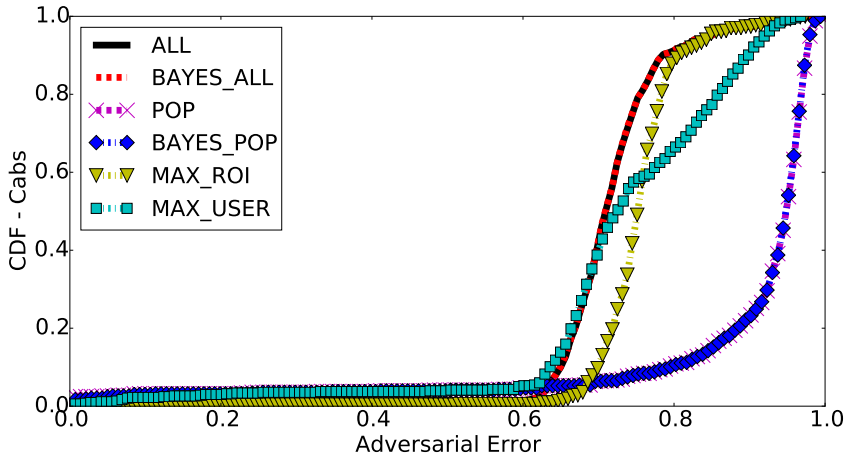

SFC

User Localization

Prior Knowledge: Location frequency, for time instances of a week

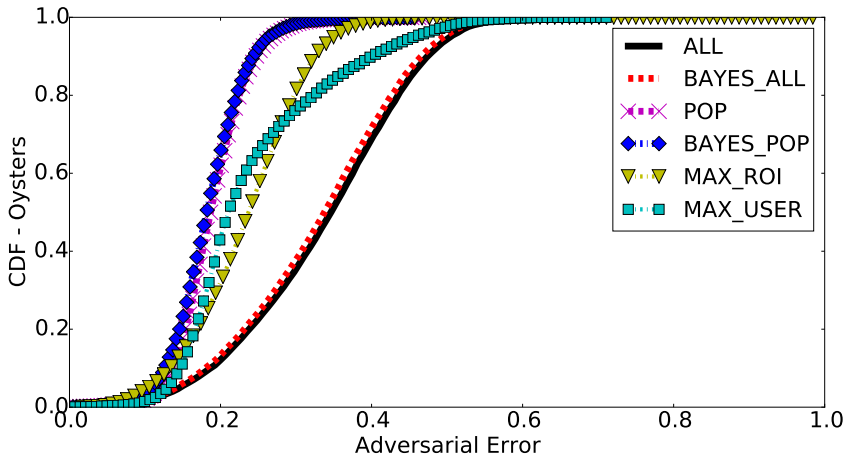

TFL

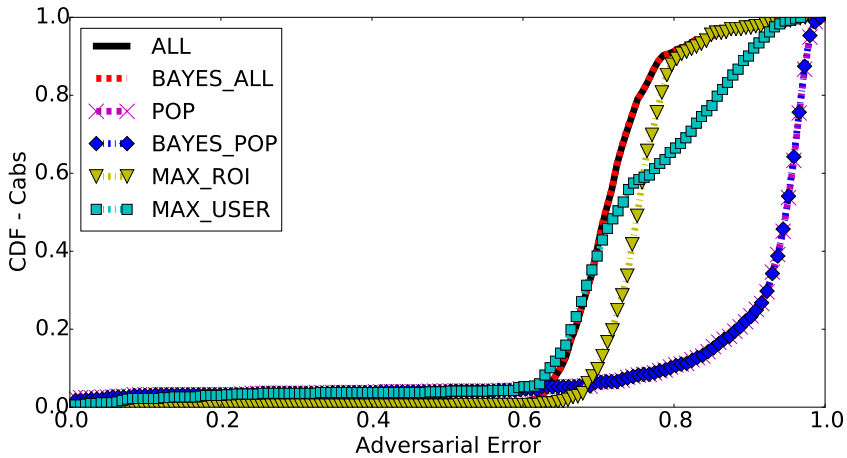

SFC

TFL vs SFC: Commuters are best localized via their most popular ROIs, whereas cabs via their last hour's ROIs

Take Aways

Take Aways

Aggregates do help the adversary to profile and localize users

Take Aways

Aggregates do help the adversary to profile and localize users

- Degree of privacy loss depends on the prior

Take Aways

Aggregates do help the adversary to profile and localize users

- Degree of privacy loss depends on the prior
- Assignment priors yield smaller privacy leakage compared to probabilistic ones

Table of Contents

Evaluation of Defenses

Evaluation of Defenses

- Location aggregates is a suitable setting for differential privacy (DP)

Evaluation of Defenses

- Location aggregates is a suitable setting for differential privacy (DP)
- How much privacy does DP provide? (with respect to ϵ and utility)

Evaluation of Defenses

- Location aggregates is a suitable setting for differential privacy (DP)
- How much privacy does DP provide? (with respect to ϵ and utility)
- **Privacy Gain (PG):** Normalized increase in adversarial error given DP aggregates compared to that with raw aggregates

Evaluation of Defenses

- Location aggregates is a suitable setting for differential privacy (DP)
- How much privacy does DP provide? (with respect to ϵ and utility)
- **Privacy Gain (PG):** Normalized increase in adversarial error given DP aggregates compared to that with raw aggregates
- Utility: Mean Relative Error (MRE)

Output Perturbation (1 / 2)

Output Perturbation (1 / 2)

Simple Counter Mechanism (SCM)

Output Perturbation (1 / 2)

Simple Counter Mechanism (SCM)

- Provides *event*-level privacy, i.e., it protects whether or not a user was in a specific location at a specific time

Output Perturbation (1 / 2)

Simple Counter Mechanism (SCM)

- Provides *event*-level privacy, i.e., it protects whether or not a user was in a specific location at a specific time
- Can be configured to achieve stronger guarantees, (e.g., ϵ-DP)

Output Perturbation (1 / 2)

Simple Counter Mechanism (SCM)

- Provides *event*-level privacy, i.e., it protects whether or not a user was in a specific location at a specific time
- Can be configured to achieve stronger guarantees, (e.g., ϵ-DP)

Fourier Perturbation Algorithm (FPA)

Output Perturbation (1 / 2)

Simple Counter Mechanism (SCM)

- Provides *event*-level privacy, i.e., it protects whether or not a user was in a specific location at a specific time
- Can be configured to achieve stronger guarantees, (e.g., ϵ-DP)

Fourier Perturbation Algorithm (FPA)

- Improves the privacy/utility trade-off by reducing the amount of noise required

Output Perturbation (1 / 2)

Simple Counter Mechanism (SCM)

- Provides *event*-level privacy, i.e., it protects whether or not a user was in a specific location at a specific time
- Can be configured to achieve stronger guarantees, (e.g., ϵ-DP)

Fourier Perturbation Algorithm (FPA)

- Improves the privacy/utility trade-off by reducing the amount of noise required
- Noise addition is done on the compressed domain (DFT), ϵ-DP per time-series

Output Perturbation (2 / 2)

Output Perturbation (2 / 2)

Privacy:

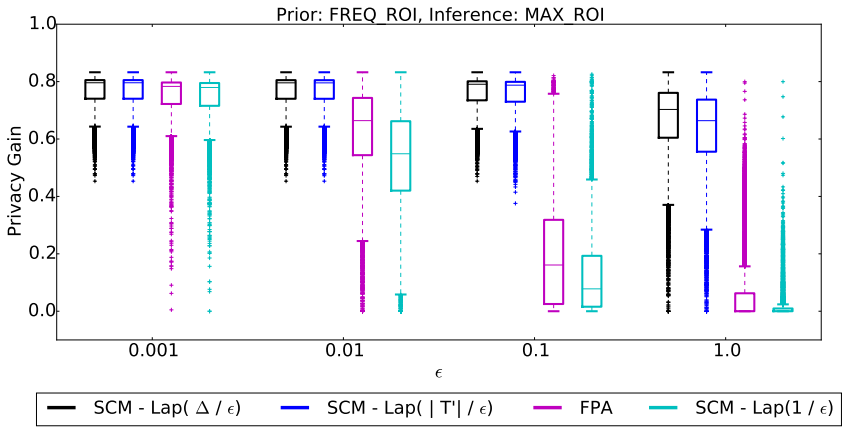

Output Perturbation (2 / 2)

Privacy:

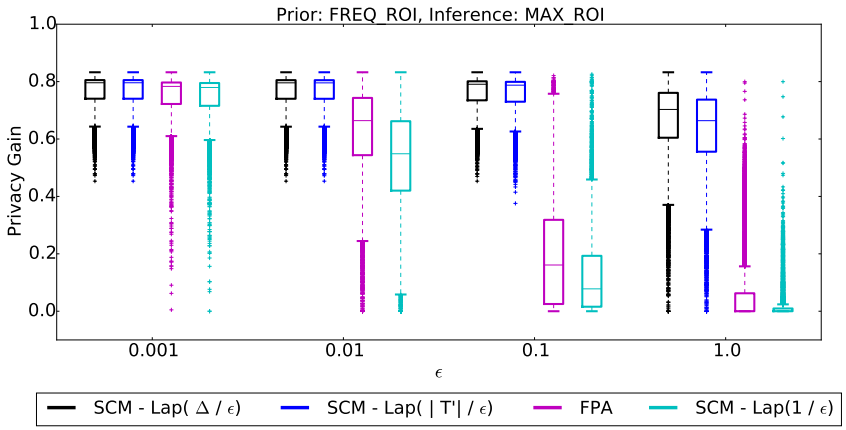

Utility:

ϵ	0.001	0.01	0.1	1.0
SCM - Lap($\lvert S\rvert \cdot \lvert T'\rvert/\epsilon$)	739.9	743.2	735.8	709.4
SCM - Lap(Δ/ϵ)	720.1	605.1	168.9	16.7
SCM - Lap($\lvert T'\rvert/\epsilon$)	719.8	549.6	123.5	12.8
FPA	117.1	11.7	1.3	0.3
SCM - Lap($1/\epsilon$)	74.4	7.8	0.9	0.1

Table 3. TFL: MRE (Utility) of output perturbation mechanisms.

Input Perturbation (1 / 2)

Input Perturbation (1 / 2)

Randomized Response

Input Perturbation (1 / 2)

Randomized Response

- SpotMe mechanism is focused on aggregate location time-series

Input Perturbation (1 / 2)

Randomized Response

- SpotMe mechanism is focused on aggregate location time-series
- Users report to be in a location with some probability p, or report the truth with probability $1 - p$

Input Perturbation (1 / 2)

Randomized Response

- SpotMe mechanism is focused on aggregate location time-series
- Users report to be in a location with some probability p, or report the truth with probability $1 - p$
- Aggregator collects user perturbed inputs and *estimates* the aggregates

Input Perturbation (2 / 2)

Input Perturbation (2 / 2)

Privacy:

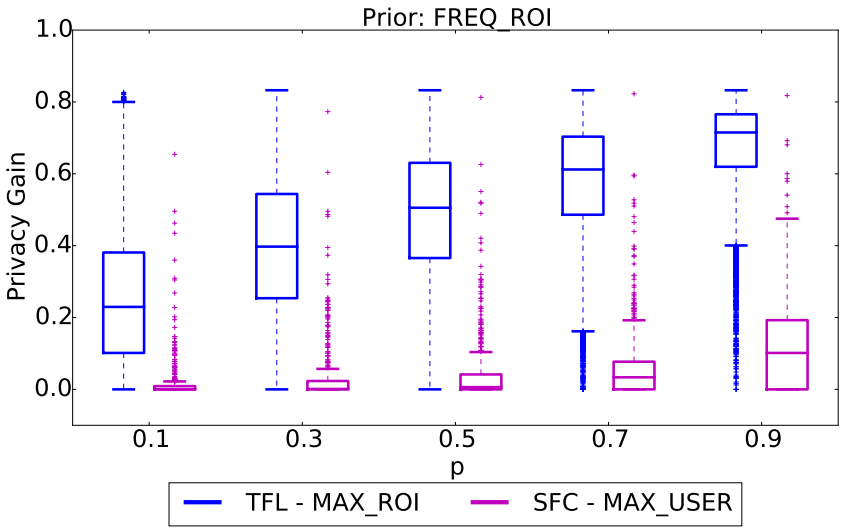

Privacy:

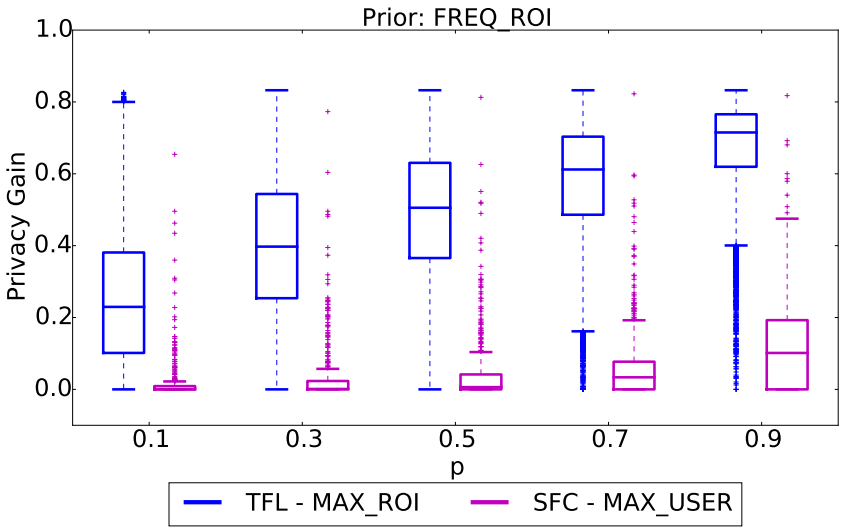

Utility:

p	0.1	0.3	0.5	0.7	0.9
TFL - MRE	2.1	3.9	6.1	9.3	17.6
SFC - MRE	0.4	0.7	1.1	1.6	2.9

Table 5. SpotMe [38]: MRE (Utility) for increasing values of p, on TFL and SFC datasets.

Table of Contents

Conclusion

Conclusion

- We presented a framework that allows us to reason about individuals' privacy loss from aggregate location time-series release

Conclusion

- We presented a framework that allows us to reason about individuals' privacy loss from aggregate location time-series release
- Location aggregates enable an adversary with some prior knowledge to profile and localize users

Conclusion

- We presented a framework that allows us to reason about individuals' privacy loss from aggregate location time-series release
- Location aggregates enable an adversary with some prior knowledge to profile and localize users
- DP mechanisms improve privacy when the utility of the time-series is poor

Conclusion

- We presented a framework that allows us to reason about individuals' privacy loss from aggregate location time-series release
- Location aggregates enable an adversary with some prior knowledge to profile and localize users
- DP mechanisms improve privacy when the utility of the time-series is poor
- Need for novel defense mechanisms for privacy-friendly mobility analytics

The end...

The end...

Thanks for your attention! Any questions?

The end...

Thanks for your attention! Any questions?

Contact Details: **apostolos.pyrgelis.14@ucl.ac.uk**