Building and Evaluating Privacy-Friendly Mobility Analytics on Aggregate Location Data

Apostolos Pyrgelis
University College London (UCL)

June 29, 2018
EPFL

About Me

Originally from Athens, Greece

Diploma & MSc from Computer Engineering and Informatics Department (CEID), University of Patras

Computer Technology Institute & Press (CTI) - EU FP7 ABC4Trust

Currently: PhD candidate at University College London (UCL) under the supervision of Emiliano De Cristofaro

Introduction

Mobility analytics are useful in modern cities for journey planning, etc.

Large-scale collection of individual users' location data prompts privacy concerns

Pseudonymization / anonymization of location traces is **ineffective**

Let There be Aggregation

Analysts are given access to aggregate location statistics, e.g., time-series

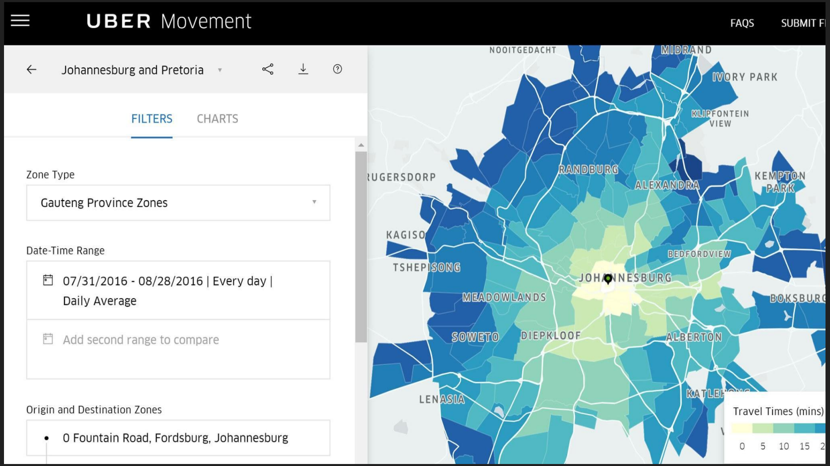

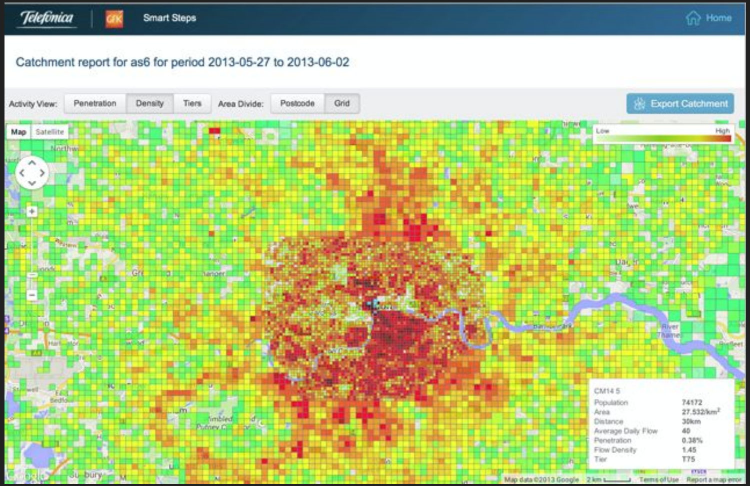

"Privacy-Friendly": User data is hidden in the crowd!

Outline

Mobility Analytics on Aggregate Location Data

Information Leakage from Aggregate Locations

- User Profiling / Localization
- Membership Inference

Future Research Directions

Mobility Analytics on Aggregate Location Data*

* Privacy-Friendly Mobility Analytics Using Aggregate Location Data. Apostolos Pyrgelis, Emiliano De Cristofaro, and Gordon Ross. ACM SIGSPATIAL 2016.

Transport for London (TFL)

Oyster Card trips including Underground, Overground, National Rail, Docklands Light Railway

Monday, March 1 to Sunday March 28, 2010 (4 weeks)

60M trips / 4M users / 582 stations

Sparse / Regular

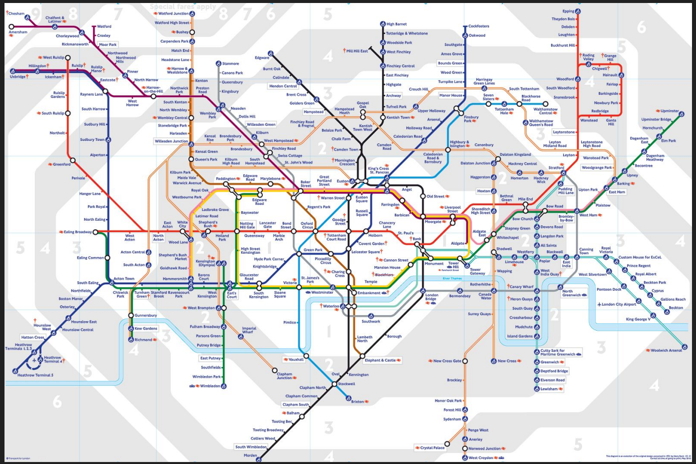

San Francisco Cab Network (SFC)

GPS mobility traces of taxis in SF

May 19 - June 8, 2008 (3 weeks)

11M coordinates / 536 cabs / 100x100 grid → 10K Regions Of Interest (ROIs) of 0.19x0.14 mi^2

Dense / Irregular

Seasonal Effects

Build hourly time-series, # of users in a ROI

TFL

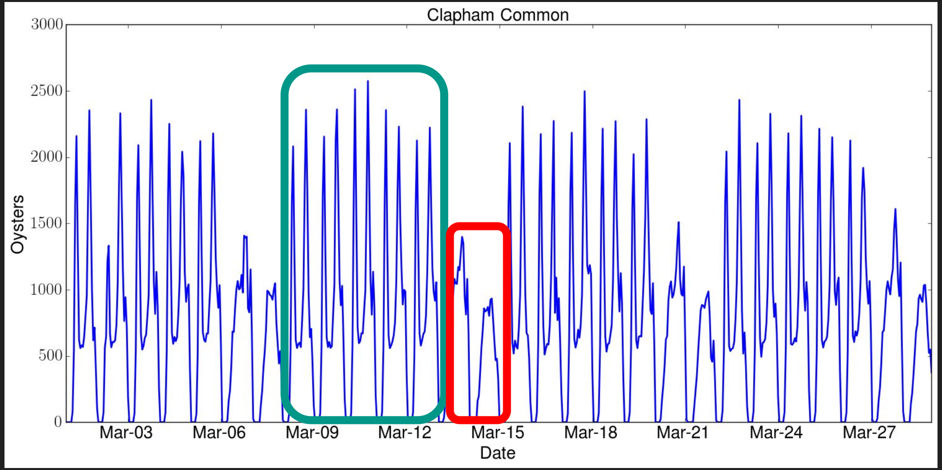

SFC

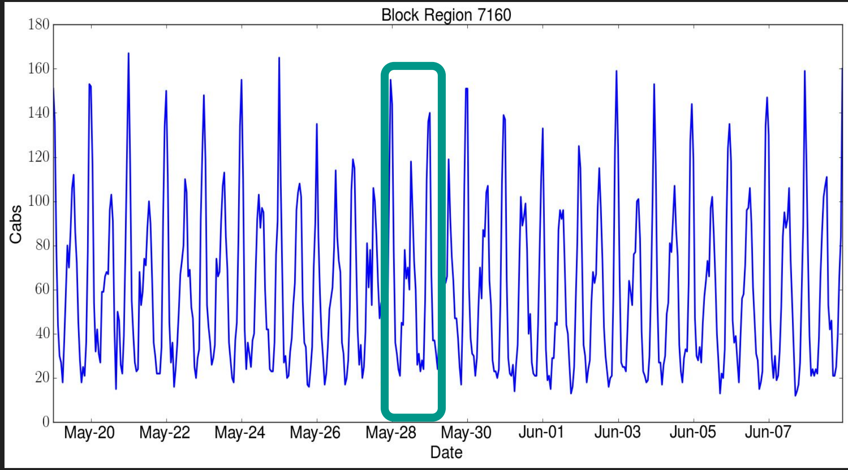

We observe daily / weekly seasonal patterns and stationarity!

Forecasting Traffic Volumes in ROIs

ARMA - Seasonal on top 100 TFL and SFC ROIs

4 Days Training / 1 Day Testing

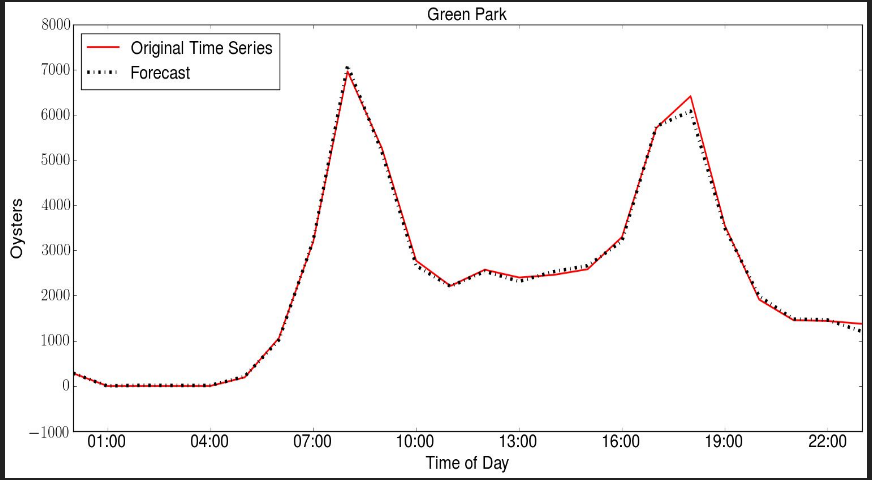

TFL, Green Park, March 25

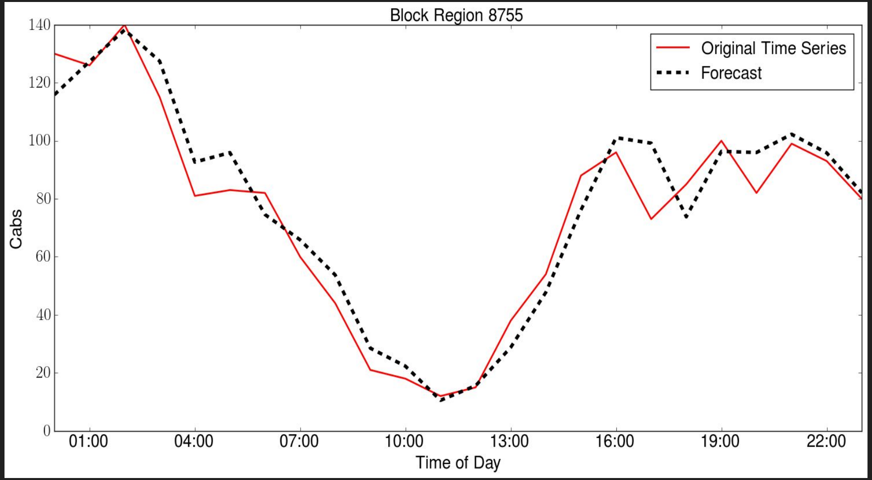

SFC, Region 8755, June 5

Seasonal Effects Improve Predictions (e.g., TFL 30x error decrease wrt. ARMA)

Detecting Mobility Anomalies

3σ rule on the forecast error / 1 Week Training, rest for Testing

TFL: 896 anomalies / SFC: 366 anomalies

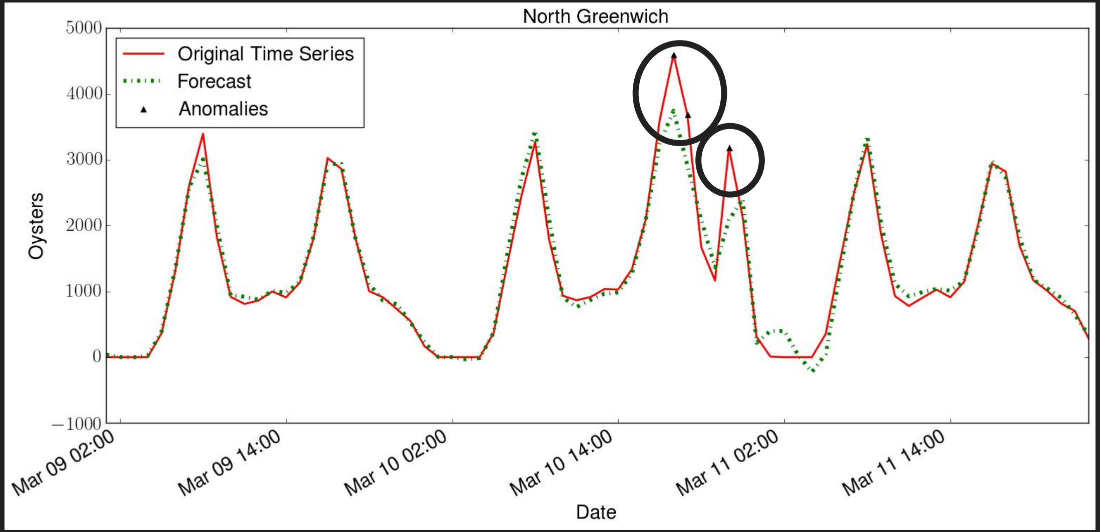

TFL, North Greenwich, March 10

Note: no ground truth!

Enhancing Predictions During Anomalies

Experiment with top 10% anomalies of TFL and SFC

VAR model with information from 10 correlated ROIs

TFL (SFC): 30% (20%) prediction improvement

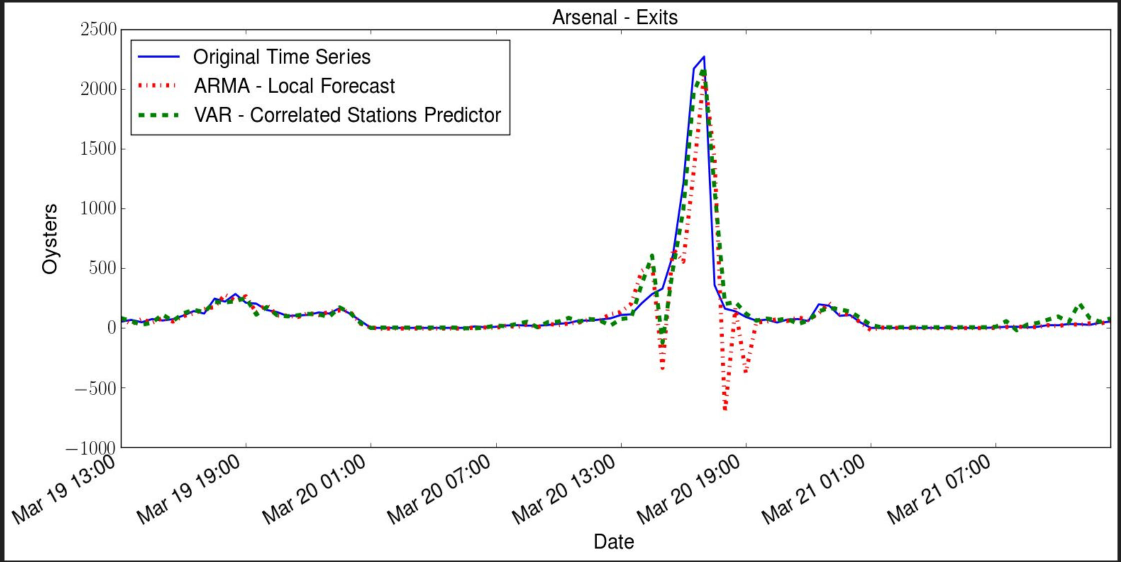

TFL, Arsenal Exit Predictions, March 20

Mobility Data Donors (MDD)*

Client Side:

Users install MDD app

MDD collects GPS coordinates

MDD encrypts the location matrix

Server Side:

Aggregator collects encrypted matrices and decrypts **ONLY** aggregate location counts - combines aggregates from multiple groups

TFL matrix (582, 2) - 4 sec encryption

10.7KB public keys, 4.54KB encrypted ROI matrix

826mJ encryption, 609mJ / 322mJ download / upload (Wi-Fi)

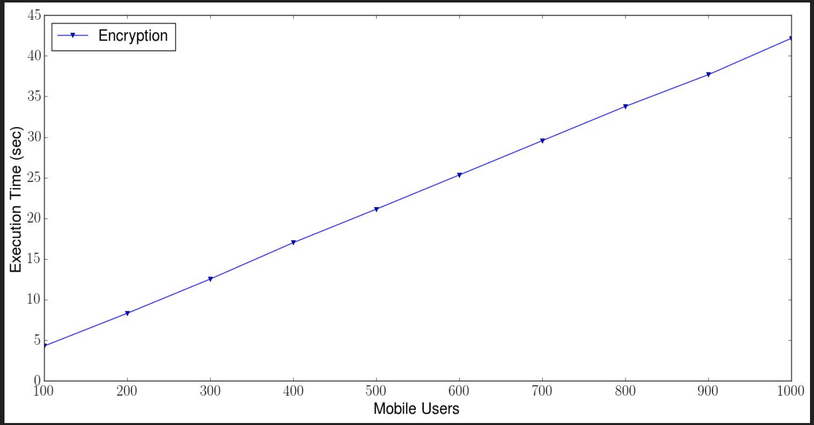

*** Efficient Private Statistics with Succinct Sketches. Luca Melis, George Danezis and Emiliano De Cristofaro. NDSS 2016.**

How Much Information Do Aggregate Locations Leak about Individual Users?*

* What Does the Crowd Say About You? Evaluating Aggregation-Based Location Privacy. Apostolos Pyrgelis, Carmela Troncoso, and Emiliano De Cristofaro. PoPETS 2017.

Framework Overview

An adversary with some prior knowledge about the whereabouts of individuals

Attempts to improve her knowledge based on the aggregate location time-series

Adversarial Goals:

1) Profiling: Infer the probability of a user being in a ROI at a certain time (JS-divergence)

2) Localization: Predict where the user will be (F1 score)

Privacy Loss:

Normalized reduction in adversarial error with vs without the aggregates

Adversarial Prior Knowledge

It can originate from social networks, data leaks, released location traces by providers or even personal knowledge, e.g., home / work locations

We build it over an **observation** period!

Probabilistic:

Frequency of Locations (over time)

Location Seasonality (day / week)

Assignment:

Most Popular Locations

All prior locations

Last Season (hour / day / week)

Inference Strategies

Given the prior knowledge and the aggregates over the **inference** period

1) Bayesian Update:

Posterior probability of a user being in a location at a certain time

2) MAX-ROI:

Assign the most probable users to each location, until the aggregates are consumed

3) MAX-USER:

Assign each user to her most likely location, until the aggregates are consumed

User Profiling

Observation Period: 3 (2) weeks -- **Inference Period:** 1 week TFL (SFC)
Prior Knowledge: Location Frequency, over the observation period

TFL

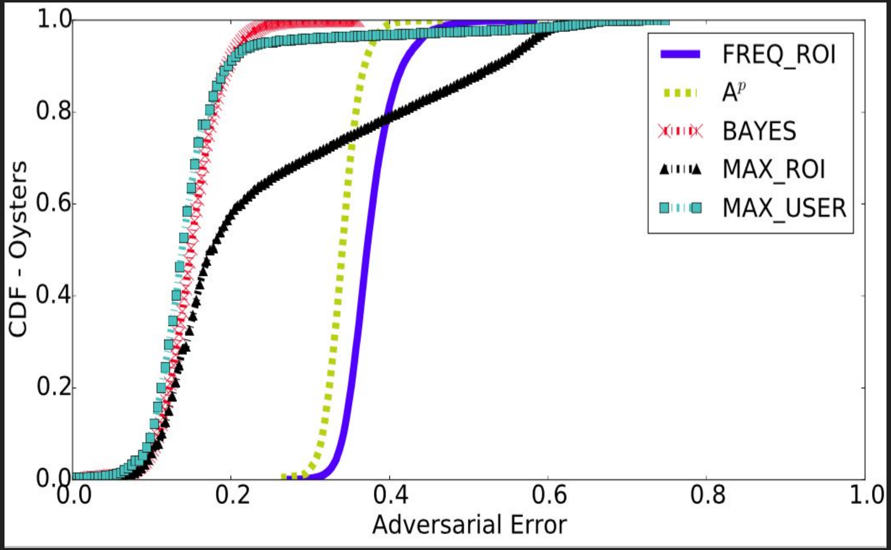

SFC

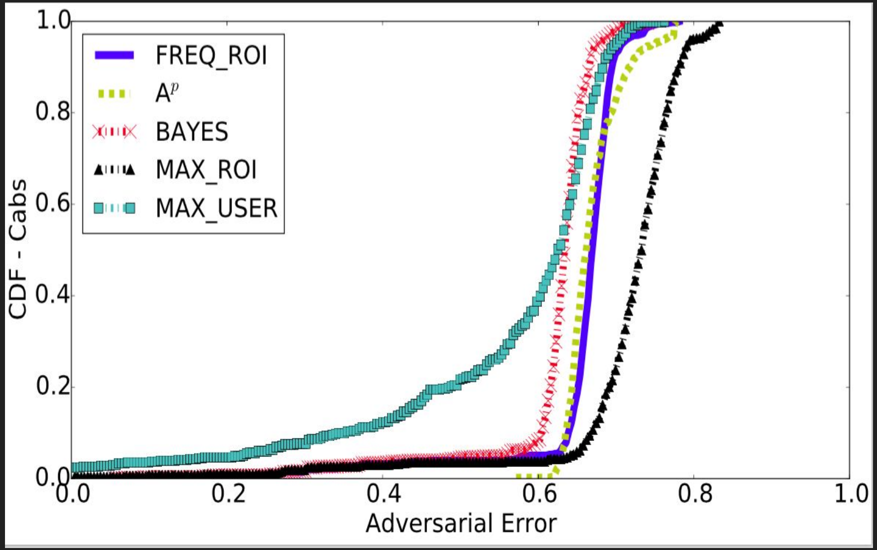

TFL vs SFC: Inferring mobility profiles from commuters is easier than those of cabs

Implications of Regular Mobility Patterns

Prior Knowledge: Location Frequency, for the time instances of any day

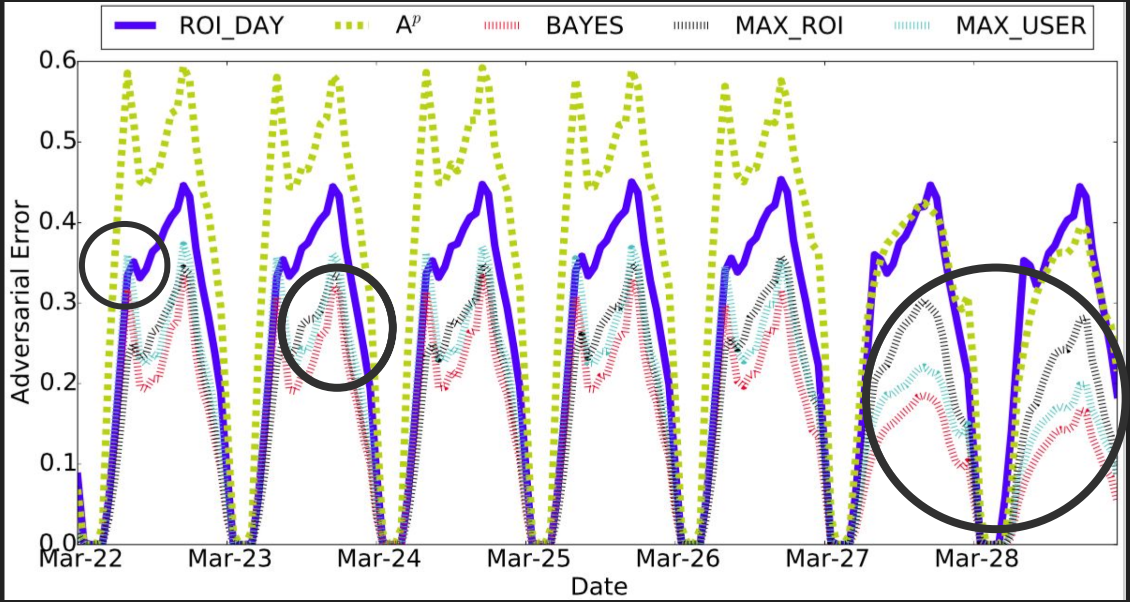

TFL

User Localization

Prior Knowledge: Location Frequency, for the time instances of a week

TFL

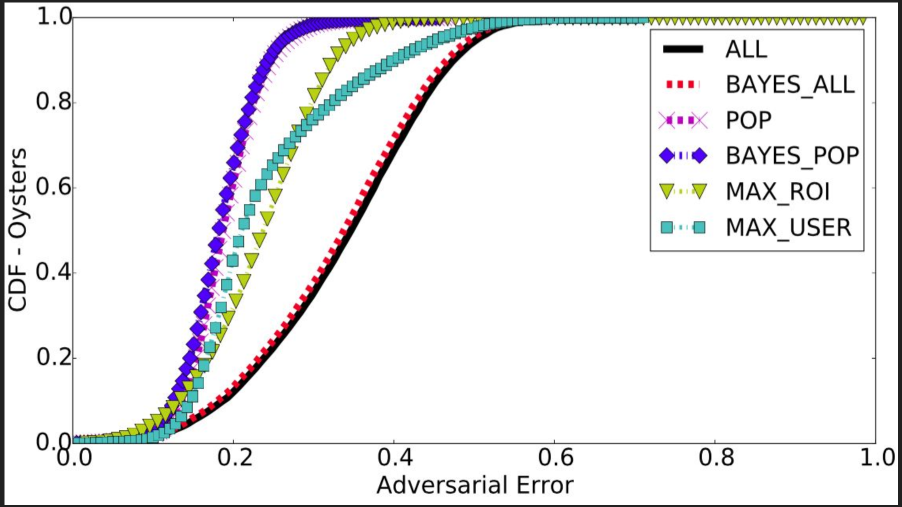

SFC

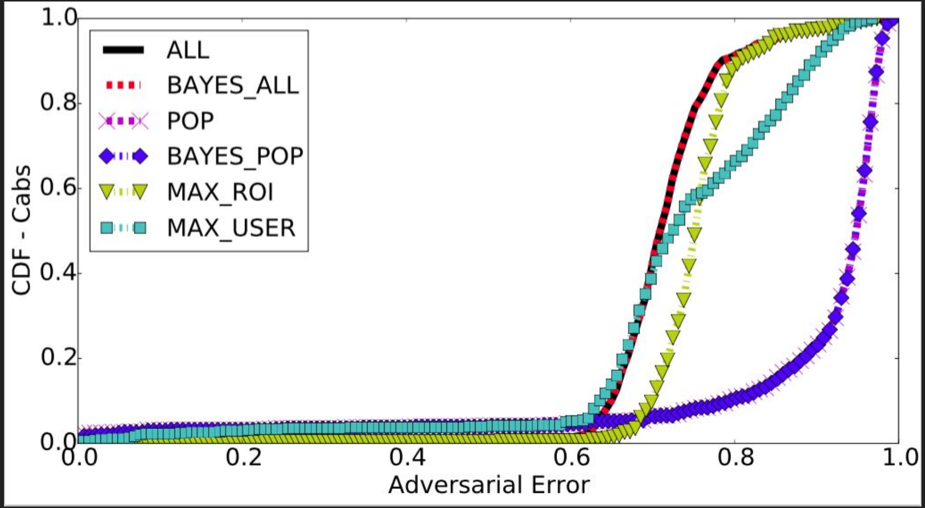

TFL vs SFC: Commuters are best localized via their most popular ROIs, whereas cabs via their last hour's ROIs

Defenses?

Aggregates do help an adversary with some prior knowledge, to profile or localize users

We can use our framework to evaluate potential countermeasures!

Privacy Gain: Normalized increase in adversarial error given the perturbed aggregates vs. raw aggregates

Utility: Mean Relative Error (MRE)

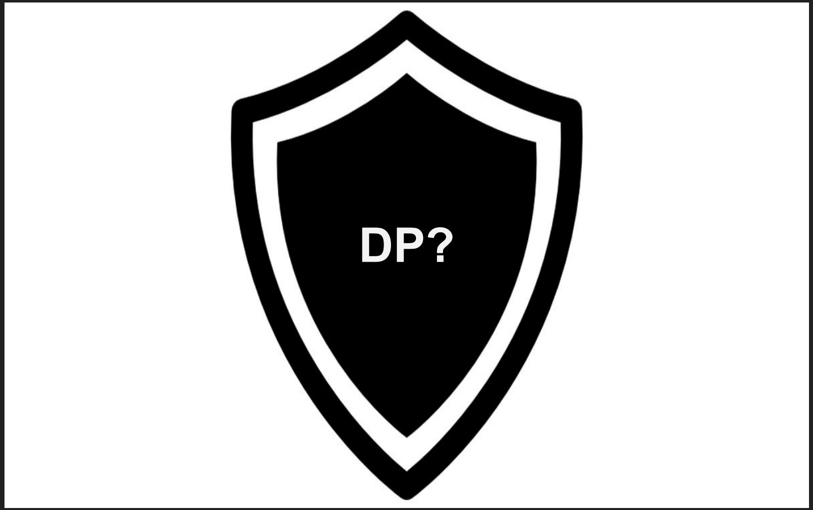

Input Perturbation

SpotMe* mechanism focused on aggregate location time-series

Users report to be in a location with probability p, or report the truth with 1-p

Aggregator collects user perturbed inputs and *estimates* the aggregates

Utility:

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

Privacy Gain:

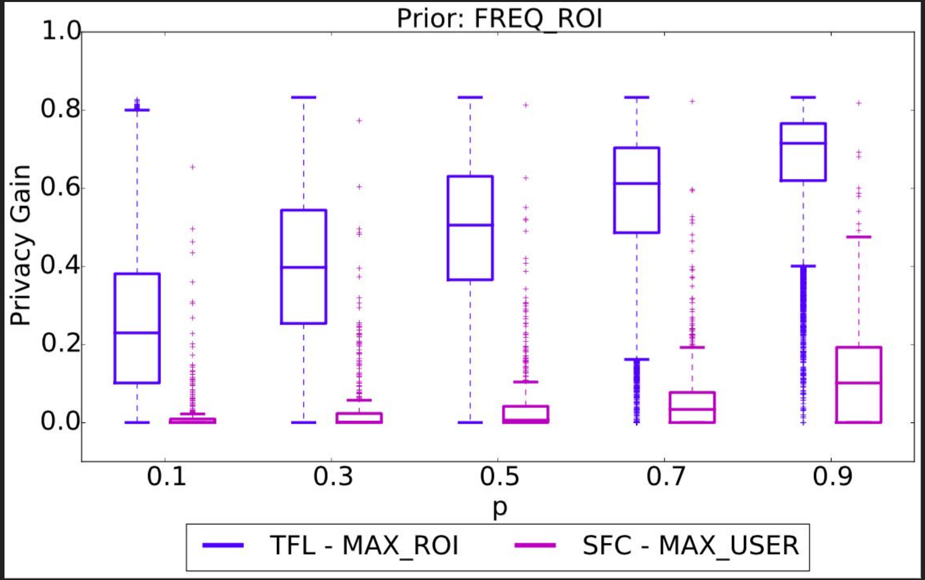

* Spot me if you can: Randomized responses for location obfuscation on mobile phones. Quercia, et al. IEEE ICDCS, 2011.

What About Membership?*

* Knock Knock, Who’s There? Membership Inference on Aggregate Location Data. Apostolos Pyrgelis, Carmela Troncoso, and Emiliano De Cristofaro. NDSS 2018. Distinguished Paper Award.

Why Membership?

Membership inference is a first step to other types of attacks, e.g., **profiling** or **localization**

Aggregates might relate to a group sharing a sensitive characteristic

Regulators can verify possibly misuse of the data

Distinguishability Game (DG)

Adversary aims at distinguishing location aggregates that include a target user from those that do not

Membership inference is a binary classification task

Supervised learning on data of the adversarial prior knowledge

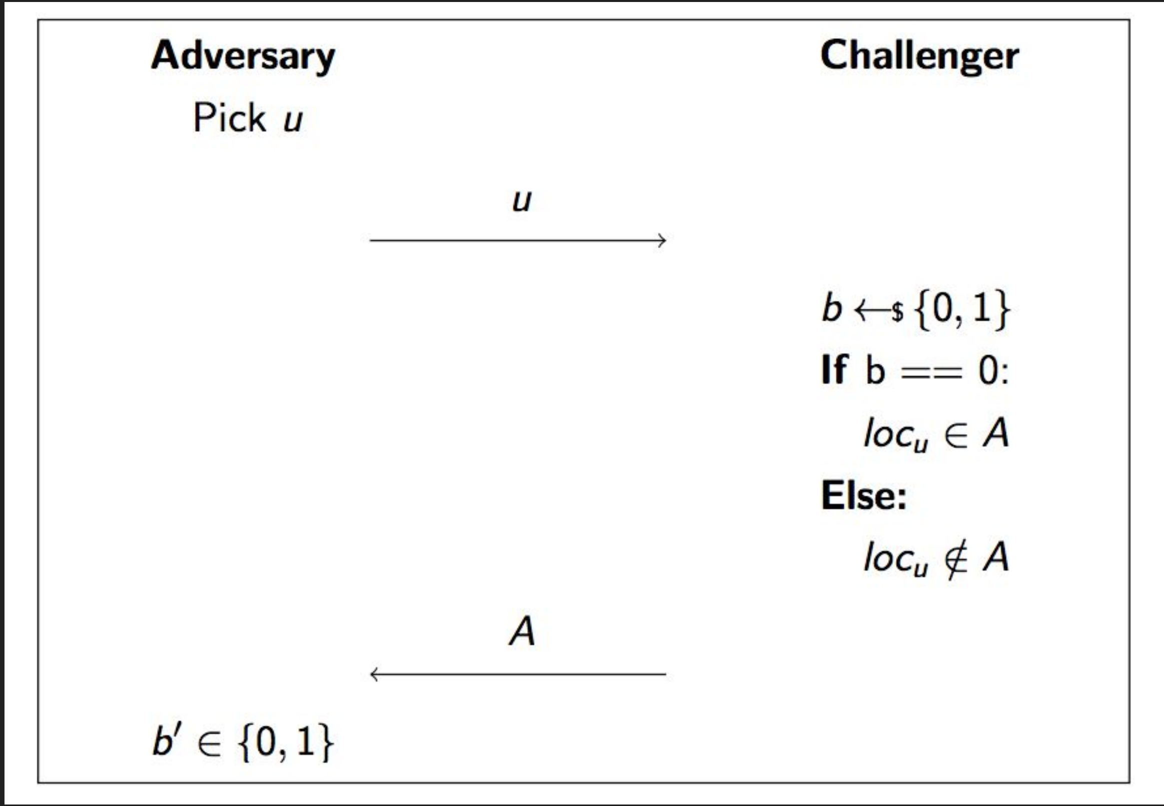

Adversarial Prior Knowledge

1) Subset of Locations: Adv knows the real locations for a subset of users (incl. her target)

2) Participation in Past Groups: Adv knows the target's participation in past aggregates

E.g., continuous data release over stable / dynamic groups

Privacy Loss

For a target, we play the distinguishability game multiple times

Privacy Loss: adversary’s advantage in winning the game over a random guess

We utilize the **Area Under Curve (AUC)** score to evaluate the classifier’s overall performance

Experimental Setup

Target Users

Randomly pick 50 users from 3 mobility groups and run membership inference

Sample & Aggregate

Balanced dataset of groups that include / exclude the target and aggregate their locations

Feature Extraction

Calculate statistics from the time-series per ROI
i.e., mean, variance, std, median, min, max, sum

Classification

Train / test a classifier:
LR, k-NN, RF, MLP

Results (Some)

TFL

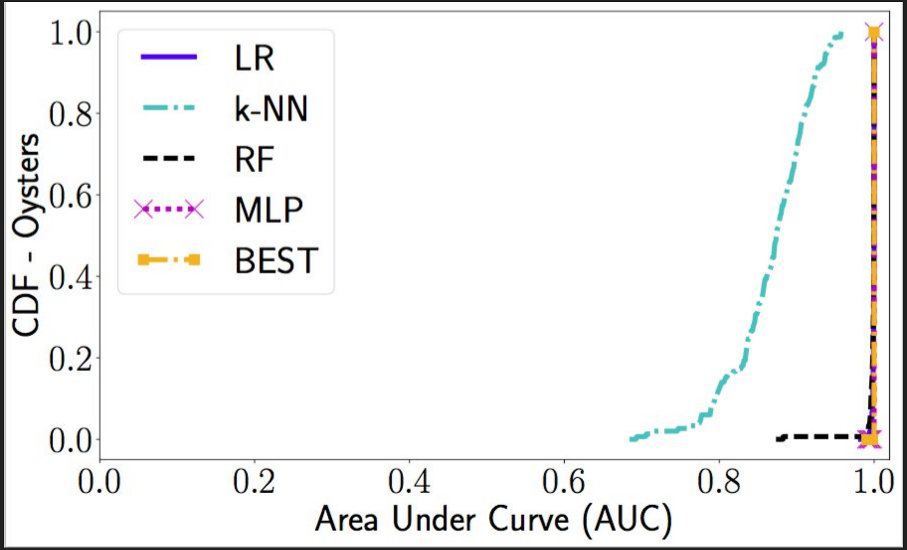

Prior: Same Groups As Released

Group Size: 1000

Inference Period: 1 week

SFC

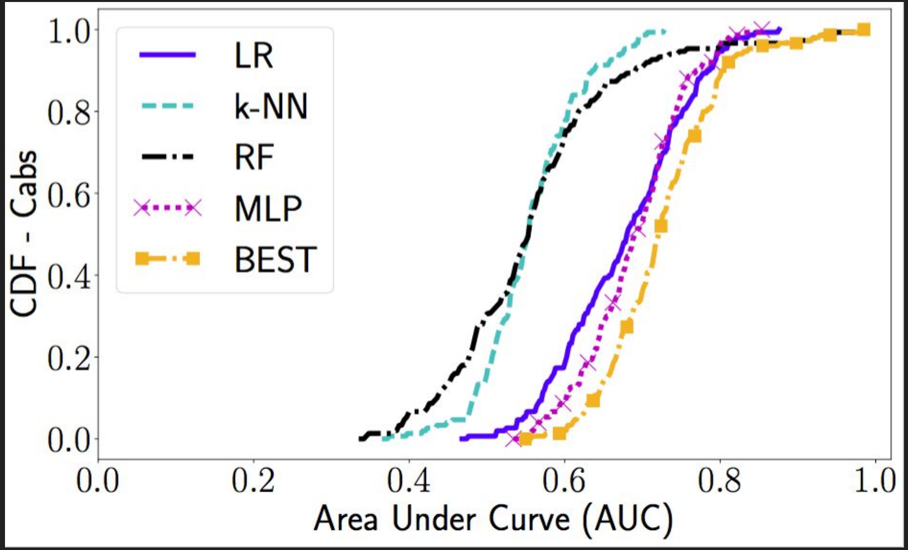

Prior: Subset of Locations

Group Size: 100

Inference Period: 1 week

Parameters of DG

Group Size

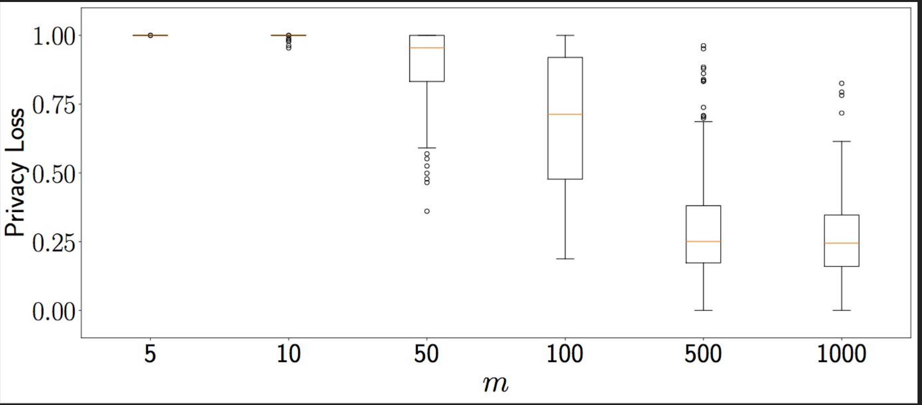

Dataset: TFL

Prior: Subset of Locations

Inference Period: 1 week

Inference Period

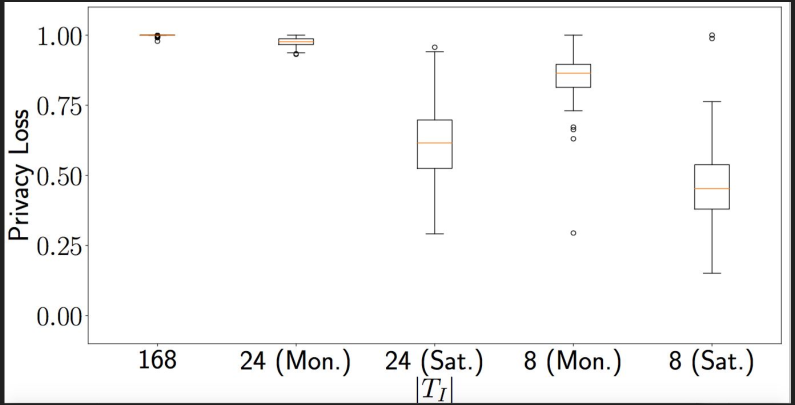

Dataset: TFL

Prior: Same Groups As Released

Group Size: 1000

Defenses?

We choose a worst-case adversary (AUC score 1.0 on *raw* aggregates)

The challenger applies a DP mechanism before sending her challenge to the adversary

LPA, GSM, FPA, EFPAG

Privacy Gain: Normalized decrease in the adversarial performance given the perturbed aggregates vs. raw aggregates

Utility: Mean Relative Error (MRE)

Results - TFL - Group Size 9.5K

Utility:

ϵ	0.01	0.1	1.0	10.0
LPA(Δ/ϵ)	3056.1	812.6	81.7	8.2
GSM	753.2	75.8	7.4	0.75
FPA	67.2	6.1	0.7	0.03
EFPAG	36.8	3.6	0.4	0.03
LPA(1/ϵ)	38.5	3.7	0.3	0.002

Privacy Gain:

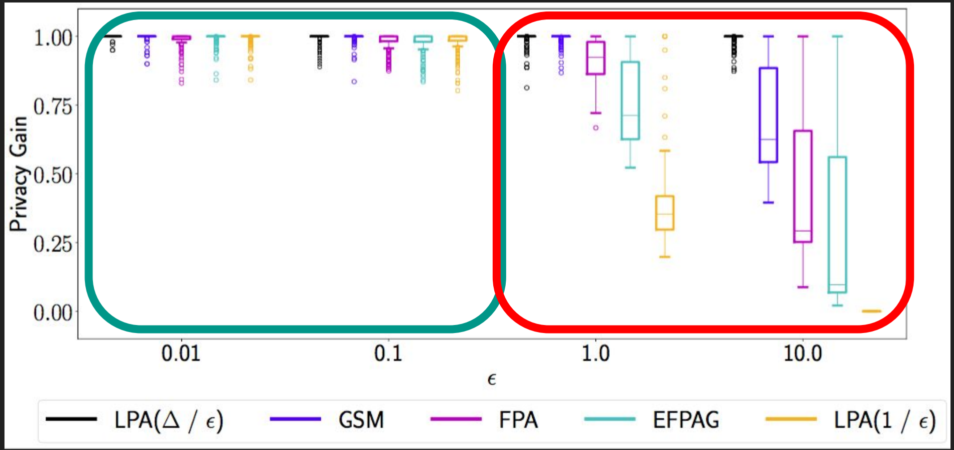

Take Aways

Aggregate location time-series are useful for mobility analytics

But, aggregates are privacy invasive:

- Leak information about individuals
- Membership inference is feasible

DP offers good protection against inferences

But, with significant reduction in the utility of the aggregates

Our methods can be used to evaluate defense approaches!

Future Research Directions

Privacy Threats on Location Data

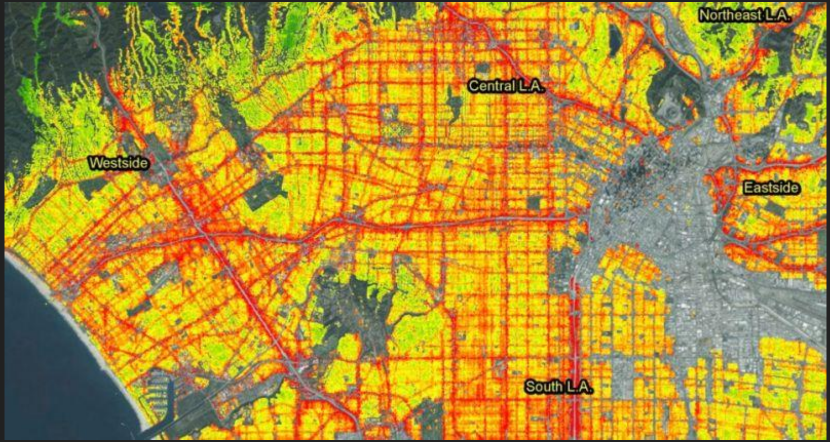

+

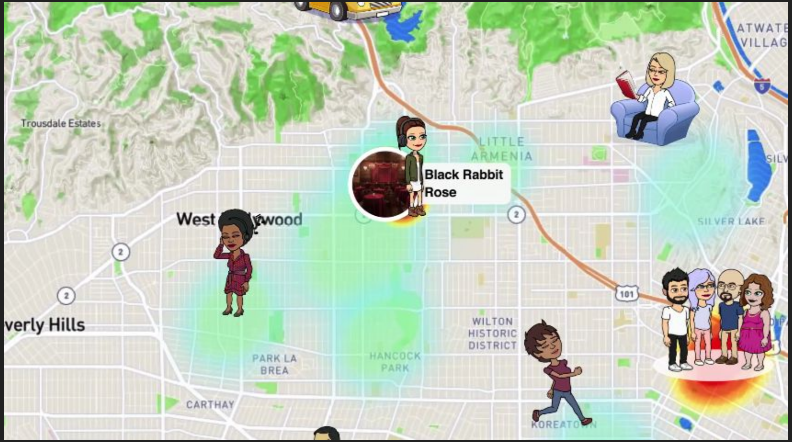

Combine aggregate location statistics with co-location information* for inference

Framework to quantify privacy leakage & evaluate protection approaches

* Quantifying interdependent privacy risks with location data. Olteanu et al., IEEE Transactions on Mobile Computing, 2017.

Health Data & Privacy Issues

Digitization of health data enables preventive medicine and research

But, create various privacy risks:

- User identification
- Discrimination
- Interdependencies

Collaboration & Contribution

Collaborative privacy-preserving data sharing frameworks*

Quantification methodologies for capturing privacy leakage

Evaluation of defense mechanisms for preventing inferences

* Unlynx: a decentralized system for privacy-conscious data sharing. Froelicher, et al. PoPETS, 2017.

Thanks for your attention!

Also, special thanks to:

Dr. Emiliano De Cristofaro

Prof. Carmela Troncoso

Dr. Gordon Ross

(Dr.) Luca Melis

Xerox University Affairs Committee grant on "Secure Collaborative Analytics".

Contact:

apostolos.pyrgelis.14@ucl.ac.uk

Web:

https://mex2meou.github.io/webpage