# Privacy preserving location based services : From centralized to federated approaches

Besma Khalfoun

N°d'ordre NNT : 2022ISAL0089

**THESE de DOCTORAT DE L'INSA LYON, membre de l'Université de Lyon**

**Ecole Doctorale** N° 512
**Mathématiques et Informatique (InfoMaths)**

**Spécialité/ discipline de doctorat : INFORMATIQUE**

Soutenue publiquement le 20/10/2022, par :
**Besma Khalfoun**

---

# Privacy Preserving Location based Services: from Centralized to Federated Approaches

---

Devant le jury composé de :

| | | |
|---|---|---|
| Musolesi, Mirco | Professeur des Universités, Université Collège London | Rapporteur |
| Nguyen, Benjamin | Professeur des Universités, INSA Val de Loire | Rapporteur |
| Goga, Oana | Chargée de recherche CNRS, Grenoble | Examinatrice |
| Lamarre, Philippe | Professeur des Universités, INSA Lyon | Examinateur |
| Ben Mokhtar, Sonia | Directrice de Recherche, CNRS | Directrice de thèse |
| Bouchenak, Sara | Professeure des Universités, INSA Lyon | Co-Directrice de thèse |

# Département FEDORA – INSA Lyon - Ecoles Doctorales

| SIGLE | ECOLE DOCTORALE | NOM ET COORDONNEES DU RESPONSABLE |
|---|---|---|
| CHIMIE | **CHIMIE DE LYON**<br>**https://www.edchimie-lyon.fr**<br>Sec. : Renée EL MELHEM<br>Bât. Blaise PASCAL, 3e étage<br>secretariat@edchimie-lyon.fr | **M. Stéphane DANIELE**<br>C2P2-CPE LYON-UMR 5265<br>Bâtiment F308, BP 2077<br>43 Boulevard du 11 novembre 1918<br>69616 Villeurbanne<br>**directeur@edchimie-lyon.fr** |
| E.E.A. | **ÉLECTRONIQUE, ÉLECTROTECHNIQUE, AUTOMATIQUE**<br>**https://edeea.universite-lyon.fr**<br>Sec. : Stéphanie CAUVIN<br>Bâtiment Direction INSA Lyon<br>Tél : 04.72.43.71.70<br>secretariat.edeea@insa-lyon.fr | **M. Philippe DELACHARTRE**<br>INSA LYON<br>Laboratoire CREATIS<br>Bâtiment Blaise Pascal, 7 avenue Jean Capelle<br>69621 Villeurbanne CEDEX<br>Tél : 04.72.43.88.63<br>**philippe.delachartre@insa-lyon.fr** |
| E2M2 | **ÉVOLUTION, ÉCOSYSTÈME, MICROBIOLOGIE, MODÉLISATION**<br>**http://e2m2.universite-lyon.fr**<br>Sec. : Sylvie ROBERJOT<br>Bât. Atrium, UCB Lyon 1<br>Tél : 04.72.44.83.62<br>secretariat.e2m2@univ-lyon1.fr | **M. Philippe NORMAND**<br>Université Claude Bernard Lyon 1<br>UMR 5557 Lab. d'Ecologie Microbienne<br>Bâtiment Mendel<br>43, boulevard du 11 Novembre 1918<br>69 622 Villeurbanne CEDEX<br>**philippe.normand@univ-lyon1.fr** |
| EDISS | **INTERDISCIPLINAIRE SCIENCES-SANTÉ**<br>**http://ediss.universite-lyon.fr**<br>Sec. : Sylvie ROBERJOT<br>Bât. Atrium, UCB Lyon 1<br>Tél : 04.72.44.83.62<br>secretariat.ediss@univ-lyon1.fr | **Mme Sylvie RICARD-BLUM**<br>Institut de Chimie et Biochimie Moléculaires et Supramoléculaires (ICBMS) - UMR 5246 CNRS - Université Lyon 1<br>Bâtiment Raulin - 2ème étage Nord<br>43 Boulevard du 11 novembre 1918<br>69622 Villeurbanne Cedex<br>Tél : +33(0)4 72 44 82 32<br>**sylvie.ricard-blum@univ-lyon1.fr** |
| INFOMATHS | **INFORMATIQUE ET MATHÉMATIQUES**<br>**http://edinfomaths.universite-lyon.fr**<br>Sec. : Renée EL MELHEM<br>Bât. Blaise PASCAL, 3e étage<br>Tél : 04.72.43.80.46<br>infomaths@univ-lyon1.fr | **M. Hamamache KHEDDOUCI**<br>Université Claude Bernard Lyon 1<br>Bât. Nautibus<br>43, Boulevard du 11 novembre 1918<br>69 622 Villeurbanne Cedex France<br>Tél : 04.72.44.83.69<br>**hamamache.kheddouci@univ-lyon1.fr** |
| Matériaux | **MATÉRIAUX DE LYON**<br>**http://ed34.universite-lyon.fr**<br>Sec. : Yann DE ORDENANA<br>Tél : 04.72.18.62.44<br>yann.de-ordenana@ec-lyon.fr | **M. Stéphane BENAYOUN**<br>Ecole Centrale de Lyon<br>Laboratoire LTDS<br>36 avenue Guy de Collongue<br>69134 Ecully CEDEX<br>Tél : 04.72.18.64.37<br>**stephane.benayoun@ec-lyon.fr** |
| MEGA | **MÉCANIQUE, ÉNERGÉTIQUE, GÉNIE CIVIL, ACOUSTIQUE**<br>**http://edmega.universite-lyon.fr**<br>Sec. : Stéphanie CAUVIN<br>Tél : 04.72.43.71.70<br>Bâtiment Direction INSA Lyon<br>mega@insa-lyon.fr | **M. Jocelyn BONJOUR**<br>INSA Lyon<br>Laboratoire CETHIL<br>Bâtiment Sadi-Carnot<br>9, rue de la Physique<br>69621 Villeurbanne CEDEX<br>**jocelyn.bonjour@insa-lyon.fr** |
| ScSo | **ScSo***<br>**https://edsciencessociales.universite-lyon.fr**<br>Sec. : Mélina FAVETON<br>INSA : J.Y. TOUSSAINT<br>Tél : 04.78.69.77.79<br>melina.faveton@univ-lyon2.fr | **M. Christian MONTES**<br>Université Lumière Lyon 2<br>86 Rue Pasteur<br>69365 Lyon CEDEX 07<br>**christian.montes@univ-lyon2.fr** |

*ScSo : Histoire, Géographie, Aménagement, Urbanisme, Archéologie, Science politique, Sociologie, Anthropologie

# Abstract

Nowadays, the proliferation of handheld devices embedded with multiple mobile sensors and the growth of fast communication and data processing technologies have contributed to the emergence of a wide variety of online services, including location-based services. These services facilitate users' daily lives with a broad range of applications that offer users personalized and customized information about their surroundings according to their location. While these services have undeniably become essential and indispensable to our society today and especially in the future, it is necessary to underline and understand the risks and threats affecting users. Indeed, large amounts of mobility data are being gathered, stored, and processed by service providers or third parties without necessarily the users' consent. As a consequence, users' privacy is threatened, and thus many sensitive information, such as the user's identity, home or workplace address, or even religious beliefs or health status can be inferred and leaked.

In this context, it becomes urgent to devise mechanisms that allow users to securely and safely access location-based services without disclosing their private lives. To address this challenge, many efforts aim to enhance privacy by proposing new location privacy protection mechanisms (LPPMs). These efforts are not only motivated by the research community, but authorities and organizations increasingly establish new laws and regulations to reframe the collection, storage, and manipulation of users' mobility data. In this direction, location privacy risk assessment (LPRA) is defined to assess the privacy risks of sharing mobility data to raise users' awareness about their privacy. In this manuscript, we use the re-identification risk, which aims at re-linking an anonymous mobility data to its originating user as a means of LPRA.

In this thesis, we first propose MOOD, a centralized user-centric protection system that aims to protect the mobility data of all users and in particular those who are not protected by any individual LPPM. MOOD uses the composition of several LPPMs and incorporates the re-identification risk assessment before publishing the protected data. However, it requires a trusted proxy server to perform both the obfuscation process and the re-identification risk assessment. Although existing protection methods aim to eliminate the trusted proxy server, the privacy risk assessment still needs to centralize the mobility data. That is why we propose

SAFER, a novel privacy risk assessment metric, developed on the user side, to estimate how unique a user's mobility data is among a group of participating users. SAFER follows a federated learning approach to build a global knowledge without accessing raw users' mobility data in a central entity. Finally, we propose EDEN, a user-side mobility data protection system that automatically selects the best LPPM and its corresponding configuration that resists the re-identification risk assessment without sending the raw mobility data outside the user's device thanks to the federated learning paradigm.

**Keywords:** privacy, location-based services, mobility data, protection mechanisms, location privacy risk assessment, re-identification attacks, uniqueness, data utility.

# Résumé

De nos jours, la prolifération des appareils mobiles embarquant de multiples capteurs et la croissance rapide des technologies de communication et de traitement de la donnée ont contribué à l'émergence d'une grande variété de services en ligne, dont les services basés sur la localisation. Ces services facilitent la vie quotidienne des utilisateurs en leur offrant des informations personnalisées et customisées sur leur environnement en fonction de leur localisation. Tout en reconnaissant qu'il est indéniable que ces services sont devenus incontournables et indispensables à notre société actuelle et surtout future, il y a lieu de souligner et d'appréhender les risques et les dangers quant à la vie privée des utilisateurs. En effet, de grandes quantités de données de mobilité sont collectées, stockées et traitées par des fournisseurs de services ou des tiers, sans forcément respecter le cadre consenti par les utilisateurs. Par conséquent, la vie privée de ces derniers est menacée et donc plusieurs informations sensibles telles que l'identité de l'utilisateur, son lieu de domicile ou de travail ou même ses croyances religieuses ou son état de santé peuvent être inférées de ces données.

Dans ce contexte, il devient urgent de concevoir des mécanismes de protection qui permettent aux utilisateurs d'accéder en toute sécurité aux services basés sur la localisation sans la crainte de dévoiler leur intimité. Pour relever ce défi, de nombreux efforts visent à développer des mécanismes de protection appelés "Location Privacy Protection Mechanisms (LPPM)". Ces efforts ne sont pas seulement motivés par la communauté scientifique mais sont de plus en plus imposés par les autorités et les pouvoirs publics en établissant de nouvelles règles et lois pour recadrer la collecte, le stockage et la manipulation de ces données. Dans ce sens, l'évaluation des risques liés à la confidentialité de la mobilité appelée "Location Privacy Risk Assessment (LPRA)" est définie afin de sensibiliser les utilisateurs aux risques engendrés par le partage de leurs données de mobilité. Dans le cas de notre étude, cette notion se traduit par l'évaluation du risque de ré-identification, c'est-à-dire le risque de réassocier une donnée de mobilité anonyme à son utilisateur d'origine.

Dans ce cadre, nous proposons tout d'abord MOOD, un système de protection centralisé centré sur l'utilisateur qui a pour but de protéger les données de mobilité de tous les utilisateurs et, en particulier, les utilisateurs orphelins qui ne sont protégés

par aucun LPPM individuel. MOOD utilise la composition de plusieurs LPPM et intègre l'évaluation du risque de ré-identification avant de publier les données protégées. Cependant, il requiert un *serveur proxy de confiance* pour procéder à la protection et à l'évaluation du risque de ré-identification. Bien que les méthodes de protection actuelles tendent à éliminer ce serveur proxy de confiance, l'évaluation du risque d'atteinte à la vie privée a toujours besoin de centraliser les données de mobilité. Pour cette raison, nous proposons SAFER, une nouvelle mesure d'évaluation du risque de confidentialité, développée du côté utilisateur pour estimer le risque de confidentialité en utilisant l'unicité des données de mobilité appelée "uniqueness". SAFER suit une approche basée sur l'apprentissage fédéré pour construire une connaissance globale sans avoir accès aux données brutes des utilisateurs de façon centralisée. Enfin, nous proposons EDEN, un système de protection des données de mobilité, développé du côté utilisateur. Il sélectionne automatiquement le meilleur LPPM et sa configuration correspondante sans envoyer les données de mobilité brutes en dehors du dispositif de l'utilisateur grâce au paradigme de l'apprentissage fédéré.

**Mots-clés:** vie privée, services basés sur la localisation, donnée de mobilité, mécanismes de protection, évaluation du risque de ré-identification, unicité de la mobilité, utilité de la donnée.

# Contents

# Chapter 1

# Introduction

## Contents

## 1.1 Context: Widespread Adoption of Handheld Devices and Location-based Services

In the last decade, information and communication technologies have known a widespread development. Most people are now equipped with handheld devices (*e.g.*, smartphones, tablets, smartwatches) embedded with high-precision mobile sensors (*e.g.*, GPS chip) to interact with their environment. According to the french institute for statistics and economic studies (INSEE), 94% of the young population (between 15-29 years old) was possessing a smartphone in 2021[1]. In addition, 75.44 billion connected devices are registered as a forecast for 2025[2]. The usage of connected devices has drastically contributed to the wide use of location-based services (later abbreviated LBSs). The latter offers users contextual and personalized information about their environment according to their location. They have changed users' daily life with a broad range of applications. Users are now able to get the optimal direction to any destination in real-time [12], forecast the weather for tomorrow or next week [120], discover nearby friends in social networks, track their physical fitness [121], play geosocial games [178], or use dating applications to meet people in their vicinity [103]. Users can also participate in crowd sensing campaigns where measurements related to their environment are linked to their locations [151, 46].

All these applications are based on the location of users which is an endless source of information. Indeed, it represents a valuable resource for urban planners, business marketers, and researchers as it can be used for traffic or health monitoring [167, 18], targeted advertising [28] or for research purposes [111]. It feeds companies such as Meta, Twitter, and Instagram, to better serve their clients and offer customized services. These companies have made and still make a dizzying profit in the business line of LBSs. It constitutes half of the global economic impact which is estimated to be 400 billion dollars in 2016 for the geospatial industry[3]. Thus, the market linked to this gold mine is huge and promising. However, the large amounts of locations gathered and stored intentionally (or not) by LBSs or any entity that may have access to the collected data constitute a real privacy threat to users.

[1] https://www.insee.fr/en/statistiques/6047983

[2] https://www.statista.com/statistics/471264/iot-number-of-connected-devices-worldwide/

[3] https://alphabeta.com/wp-content/uploads/2017/09/GeoSpatial-Report_Sept-2017.pdf

## 1.2 Privacy Threats on Mobility Data

The extensive usage of LBSs has generated huge amounts of information regarding users' locations. This data is continuously and increasingly gathered, stored, and manipulated by service providers. The latter or other entities that might have access to the collected data (*e.g.*, accidentally or through an attack) may exploit it fraudulently to infer and leak sensitive information about individuals [132, 239]. For instance, mobility data may very well reveal a user's home and workplace, health status, or even religious or sexual preferences if the latter regularly visits health centers, worship places, or libertine places respectively [99, 79]. In addition, the mobility prediction threat can put the user at risk. For instance, a malicious adversary can track the user's whereabouts, build a prediction model [54, 251] and then guess the next location of the user. This information may help a potential robber who wants to break into a user's house when the latter is not at home [234]. Moreover, other curious attackers aim at discovering social relationships between users (*e.g.*, friends, coworkers *etc.*,) [35, 233] or inferring points-of-interest (POIs) and their semantics to deduce relevant personal information about users without getting their consent [259, 133].

Furthermore, the disclosure of a user's identity is also jeopardized by re-identification attacks, *i.e.*, attacks where an anonymous location data is re-associated to its originating user based on previously collected data [152, 193]. For example, the journalists from the New York Times were able to re-identify and track the whereabouts of ex-president Trump from a dataset of more than 50 billion location pings from more than 12 million users' mobile devices [225]. Also in 2019, 460 million mobility records from more than 140,000 phones and tablets were leaked to the Norwegian broadcaster NRK[4] from which many individuals were re-identified. This is because human mobility acts as a fingerprint that uniquely identifies users as demonstrated by De Montjoye *et al.* [64]. In this paper, the authors show that only four randomly selected mobility points are sufficient to re-identify more than 95% of the users in a dataset of 1.5 million users. In this context, it is becoming increasingly important to devise mechanisms to address, deal with these threats and preserve users' location privacy.

[4]https://www.bbc.com/news/technology-59063766

## 1.3 Countermeasures: Location Privacy Protection and Risk Assessment

Due to the high number of attacks and threats affecting mobility data, location privacy became one of the main concerns of authorities and organizations and has known considerable attention from users' perspectives [27]. Many privacy-enhancing technologies (known as PETs) are devoted to safeguarding users' private lives. Some of them act on legal aspects by introducing new laws and regulations to strengthen data privacy. For instance, the California Consumer Privacy Act (CCPA) in the United States and the European Union general data protection regulation (EU GDPR). The latter established a regulation to integrate privacy-by-design in the development of new applications and location-based services (Article 78 in EU GDPR regarding technical and organizational data protection measures [183]). In addition, the EU legislation also encouraged organizations collecting or processing personal data to assess privacy risks for individuals before data release (Article 35 in EU GDPR regarding risk assessment [182]). This was also recommended by the National Institute of Standard and Technology (NIST) [114] and widely enforced by the National Commission on Informatics and Liberty (CNIL) in France. It states that a data controller shall carry out a privacy impact assessment, when the processing of users' data leads to disclosing personal information about them, for example, work performance, financial situation, health status, and sensitive mobility patterns. This risk is evaluated thanks to procedures and guidelines proposed by the CNIL.

In addition to regulations, technologies provide a practical way to enforce data privacy. In this direction, many efforts in the literature aim to develop location privacy protection mechanisms (LPPMs) that act directly on data by applying a set of techniques such as anonymization, perturbation, generalization, and fake data generation. Their main goal is to protect users' location privacy against several attacks. Therefore, they are complementary to the legal framework and are likely to become mandatory for companies that process the data to protect users' rights. However, the application of such LPPMs may have an impact on the utility of the resulting data and thus on the quality of the proposed service. Specifically, the higher the distortion of obfuscated data the lower its exploitability. Therefore, the best LPPM is the one that achieves the best privacy *vs.* utility tradeoff while assessing

the privacy risks before sharing the mobility data.

## 1.4 Problem Statement

**No One-size-fits-all LPPM.** After an extensive review of the literature on LPPMs and their effectiveness, we have noticed that users' mobility data is affected differently by LPPMs. Even more, the mobility of the same user is not protected to the same degree by two different LPPMs. This is due to the sensitivity of mobility data which generally changes according to multiple factors. For instance, the semantic of visited locations, the frequency of visits, and their duration may impact the effectiveness of protection methods (*e.g.*, being at a HIV center is more sensitive than being at the supermarket). Therefore, a data privacy officer (DPO) may decide to increase the protection level by adding more noise or by generating more fake data to confuse the attacker or by simply deleting data which is subject to privacy threats. However, the former can lead to overprotecting mobility data while deteriorating its utility and the latter may engender a large data loss. Unfortunately, few existing LPPMs try to adapt their configuration parameters to achieve the privacy *vs.* utility tradeoff. The focus is generally on the majority of users, and little attention is given to the minority of users, called *orphan users* who are still subject to privacy threats despite the use of LPPMs (*e.g.*, re-identification threat). In this direction, we propose MOOD, a novel user-centric solution that aims at protecting all users' mobility data and in particular *orphan users*. MOOD splits mobility traces into multiple sub traces and protects them by combining off-the-shelf LPPMs. The selected composition of LPPMs is assessed by the resilience to the re-identification risk.

**Centralization of Location Privacy Risk Assessment and Protection.** Location privacy risk assessment (LPRA) ensures that only protected mobility data is shared. It consists of several metrics that evaluate the privacy exposure engendered by sharing the mobility data. To reach this objective, companies, and organizations that exploit mobility data generally rely on guidelines and existing procedures that lack quantitative values. That is why the research community has been actively proposing quantitative metrics which better assess the privacy risk and thus improve

the anonymization process of mobility data. In this direction, uniqueness is a well-recognized metric to quantify the privacy bound of human mobility. It measures the similarity between users' mobility in the spatial and temporal space using randomly selected data points. The main concern of such a metric is that it requires centralizing raw users' mobility data in a single entity to verify the similarities between users' mobility. It further constitutes a single point of failure that may cause serious privacy breaches and data leakage if the entity holding and processing the data is compromised. To avoid centralizing raw data for the LPRA, we propose SAFER, a federated location privacy risk assessment based on the concept of uniqueness. We rely on a classifier trained in a federated learning way to learn about users' mobility while the latter remains private on the users' devices. The classifier helps in the construction of anonymity sets whose size reflects the uniqueness of mobility data.

Although some existing LPPMs do not centralize mobility data in a trusted entity to obfuscate it, they do lack the LPRA. Therefore, we propose EDEN, a federated approach to enforce location privacy protection through location privacy risk assessment. It is a user-side protection system that performs a re-identification risk assessment in a federated learning way. The latter helps in comparing LPPMs and thus choosing automatically the one which passes the re-identification test while maintaining a good utility.

## 1.5 Summary of Contributions

This manuscript consists of three contributions that fall on two axes: location privacy protection and location privacy risk assessment. Specifically, we propose new privacy-preserving protocols for mobility data while assessing the privacy risk of sharing it. Our solutions maximize privacy while maintaining a high utility of the protected data. First, we design a fine-grained user-centric protection mechanism where the data is centralized in a trusted entity that performs re-identification as a means of LPRA. Then, we propose a federated approach to conduct LPRA using the concept of uniqueness. Finally, we present a user-side privacy protection mechanism through re-identification risk assessment following the federated learning approach. In this section, we briefly describe the three contributions of the thesis. The following chapters will go deeper into each work with an extensive experimental evaluation of

real-world mobility datasets.

**(C1) MOOD: A Centralized Location Privacy Protection**

We propose MOOD, a centralized user-centric fine-grained multi-LPPM system for data publishing. Its main objective is to protect all users' mobility data, and in particular orphan users who are vulnerable to re-identification despite the application of LPPMs. MOOD combines multiple off-the-shelf LPPMs to protect a mobility dataset in front of state-of-the-art re-identification attacks. Specifically, MOOD applies various LPPMs on a given trace and chooses the combination which better resists the re-identification risk assessment while distorting the mobility data the least. It can reach between 97.5% and 100% of effectively protected data on various real-world mobility datasets with a reasonable data utility and a low data loss.

**(C2) SAFER: a Federated Learning Approach for Location Privacy Risk Assessment**

SAFER, a novel privacy risk assessment system that allows users to determine locally, on their device, how unique their mobility data is among a group of participating users, thereby raising awareness of privacy risks associated with sharing this data. In contrast to state-of-the-art solutions, SAFER does not require centralizing the mobility data of all users in a trusted server, but rather follows a federated learning approach, which allows assessing the uniqueness of mobility data while the latter remains private on the user's premises. Specifically, to assess the uniqueness of mobility data, SAFER trains a machine learning classifier to determine how many other users hold similar mobility data and infer anonymity sets. We carry out extensive experiments on four real-world mobility datasets of different types (GPS data, call detail records). Additionally, we consider the uniqueness of entire trajectories rather than just picking random data points (*i.e.*, exact locations or points-of-interest). The results of our experiments show that SAFER can successfully quantify the uniqueness of mobility data in a distributed manner, with comparable results to that of a well-established centralized baseline. We evaluate SAFER with up to 10,000 mobile users and demonstrate its scalability. Finally, through the implementation of a state-of-the-art re-identification attack, we illustrate that the data estimated as unique by SAFER is indeed at high risk of re-identification if it falls between the hands of a malicious entity.

**(C3) EDEN: Enforcing Location Privacy Protection through Location Privacy Risk Assessment: a Federated Learning Approach**

We present EDEN, a user-side mobility data protection system for crowd sensing applications. It is the first solution that selects automatically the best LPPM and its corresponding configuration (*i.e.*, among a set of LPPMs/configurations) without sending raw mobility traces outside the user's device. We reach this objective by relying on a federated learning approach. Specifically, for a given mobility trace, EDEN applies each LPPM to the raw trace and evaluates both : (1) the re-identification risk of the trace using this LPPM thanks to a federated user re-identification attack model and (2) the corresponding data utility. The evaluation of EDEN on five real-world mobility datasets shows that EDEN outperforms state-of-the-art single LPPMs reaching a better privacy *vs.* utility tradeoff.

## 1.6 Thesis's Results

The contributions of the thesis were the basis for four publications in international conferences/journals and workshops, as well as for three publications in national conferences and workshops.

### 1.6.1 Publications

#### International Conferences and Workshops

- Besma Khalfoun, Sonia Ben Mokhtar, Sara Bouchenak, and Vlad Nitu. 2021. EDEN: Enforcing Location Privacy through Re-identification Risk Assessment: A Federated Learning Approach. Proc. ACM Interact. Mob. Wearable Ubiquitous Technol. 5, 2, Article 68 (June 2021), 25 pages.

  DOI: https://doi.org/10.1145/3463502

- Besma Khalfoun, Mohamed Maouche, Sonia Ben Mokhtar, and Sara Bouchenak. 2019. MooD: MObility Data Privacy as Orphan Disease: Experimentation and Deployment Paper. In Proceedings of the 20th International Middleware

Conference (Middleware '19). Association for Computing Machinery, New York, NY, USA, 136–148.

DOI: https://doi.org/10.1145/3361525.3361542

- Besma Khalfoun, Sonia Ben Mokhtar, Sara Bouchenak, Towards User-Side Uniqueness Assessment of Mobility Data with a Federated Learning Approach, the 16th EuroSys Doctoral Workshop (EuroDW 2022).

- Amina Ben Salem, Besma Khalfoun, Sonia Ben Mokhtar, Afra Mashhadi: poster: Quantifying Fairness of Federated Learning LPPM Models, the 20th ACM International conference on Mobile Systems, Applications, and Services (Mobisys 2022 Posters).

  DOI: https://doi.org/10.1145/3498361.3538788

### National Conferences and Workshops

Those conferences have peer reviews but no proceedings.

- Besma Khalfoun, Mohamed Maouche, Sonia Ben Mokhtar, Sara Bouchenak: MooD: MObility Data Privacy as Orphan Disease. Compas'2019.

- Yanis Meziani, Besma Khalfoun, Sara Bouchenak, Sonia Ben Mokhtar And Vlad Nitu: Enforcing Adaptive Location Privacy With Federated Learning. Compas'2020.

- Besma Khalfoun, yanis meziani, Sara Bouchenak, Sonia ben mokhtar, Vlad Nitu: Enforcing Adaptive Location Privacy with Federated Learning. The 10th Atelier sur la Protection de la Vie Privée (APVP'2020).

### Ongoing Submissions

- (Under Review) Besma Khalfoun, Sonia Ben Mokhtar, Sara Bouchenak: SAFER: User-Side Uniqueness Assessment of Mobility Data with a Federated Learning Approach. Ubicomp 2022

### 1.6.2 Developed Software

- EDEN: A toolkit that runs a federated user re-identification risk assessment (FURIA) on mobility traces and uses its results to select the appropriate LPPM. Accessible LPPMs are (Geoi, Promesse, and TRL) and accessible utility metric is area coverage.
  https://github.com/bkhalfoun/EDEN
- MOOD: A toolkit that runs the location privacy protection on a given mobility dataset for data publishing. It provides access to LPPMs (HMC, Geoi, TRL) and re-identification attacks: AP-Attack, POI-Attack and PIT-Attack.
  https://github.com/bkhalfoun/mood
- SAFER: A toolkit that computes uniqueness in FL by constructing anonymity sets whose the size reflects its uniqueness. It offers a script to run the centralized baseline system of De Montjoye [64].
  https://github.com/bkhalfoun/safer (private for the moment for double blind review)

### 1.6.3 Communications

Our contributions were presented in different national and international conferences, workshops, and winter schools. The list of communications is listed in Table 1.1.

**Table 1.1:** *List of communications*

| Event | Date | Location | Title |
|---|---|---|---|
| Eurosys Doctoral Workshop 2022 | April, 2022 | Rennes, France | Towards User-Side Uniqueness Assessment of Mobility Data with a Federated Learning Approach |
| GDR, RSD ASF Winter School | March, 2022 | Pleynet, France | SAFER: User-Side Uniqueness Assessment of Mobility Data with a Federated Learning Approach |
| Ubicomp'21 | September, 2021 | Virtual | EDEN: Enforcing Location Privacy through Re-identification Risk Assessment: A Federated Learning Approach |
| IRIXYS Workshop | December, 2021 | Lyon,France | EDEN: Enforcing Location Privacy through Re-identification Risk Assessment: A Federated Learning Approach |
| IRIXYS Workshop | January, 2021 | Virtual | EDEN: Enforcing Location Privacy through Re-identification Risk Assessment: A Federated Learning Approach |
| APVP'2021 | June, 2021 | Virtual | EDEN: Enforcing Location Privacy through Re-identification Risk Assessment: A Federated Learning Approach |
| IRIXYS Workshop | June, 2020 | Virtual | MOOD: Mobility data privacy for orphan desease |
| GDR, RSD ASF Winter School | February, 2020 | Pleynet, France | MOOD: Mobility data privacy for orphan desease |
| Middleware'19 | December, 2019 | California, USA | MOOD: Mobility data privacy for orphan desease |
| Compas'19 | June, 2019 | Biarritz, France | MOOD: Mobility data privacy for orphan desease |

## 1.7 Structure of the Manuscript

The thesis is structured as follows. First, in Chapter 2, we present a state of the art on location privacy. We first introduce mobility data with its diverse forms and the privacy threats affecting it and in particular the re-identification threat. Then we review the literature on location privacy protection and location privacy risk assessment and we finish with a background on the federated learning paradigm. In Chapter 3, we present MOOD, a centralized user-centric method that combines multiple LPPMs to protect mobility traces in a fine-grained way against re-identification attacks. Then we introduce SAFER in Chapter 4, a novel metric to conduct LPRA in a federated learning approach leveraging the concept of uniqueness. After that, we present EDEN in Chapter 5, a novel user-side privacy protection system to enforce location privacy through re-identification risk assessment following

a federated learning approach. Finally, in Chapter 6, we conclude the manuscript with a summary of the contributions and a discussion of future research directions.

# Chapter 2

# Background and Related Work

## Contents

## 2.1 Background on Mobility Data and Privacy Threats

With the pervasiveness of handheld devices and the continuous growth of communication networks, location-based services are becoming more and more prominent in users' daily lives. They provide contextual and customized information to users' requests according to their location. These services are data-greedy, which pushes users to actively and increasingly disclose their mobility data to enjoy a better service. Hence, large amounts of mobility data are gathered, stored, and manipulated, which is a double-edged sword. On one side, mobility data helps LBS providers to improve the quality of their services, but on the other side, a misuse of the data by a curious LBS provider or a malicious entity that might have access to the collected data may reveal sensitive and private insights that violate users' privacy. In this section, we define mobility data in its diverse forms and present the main threats affecting it with a particular focus on the re-identification threat.

### 2.1.1 Mobility Data in its Diverse Forms

A mobility record or a location record is a spatio-temporal point $r = (lat, lng, t)$ associated to a given user, where *lat* and *lng* respectively correspond to the latitude and the longitude of GPS coordinates, and $t$ is a timestamp. A mobility record may correspond to (i) the actual location of a user extracted thanks to the GPS sensor embedded in the user device, (ii) the location of the closest cell tower from which a user is phoning or texting *i.e.*, a call detail record (CDR) (as depicted in the left part of Figure 2.1), or (iii) the centroid of a point-of-interest (POI), *i.e.*, a place where a user stopped for a significant time (as depicted in the central part of Figure 2.1). Finally, a sequence of mobility points $\{r_1, r_2, .., r_n\}$ constitutes a mobility trace or a trajectory $T$, as illustrated in the right part of Figure 2.1. Trajectories can be limited to a given duration (*i.e.*, length), such as trajectories of 30 minutes, 1 hour, and so on.

In addition to the spatial and temporal information of mobility data (record or trajectory), there are other possible attributes, such as the speed of the user's

movement, the direction of travel, and the altitude .... In the context of crowd sensing applications [89], mobility data can be mapped to environmental measurements, such as air pollution measurements [68] and radioactivity level [46].

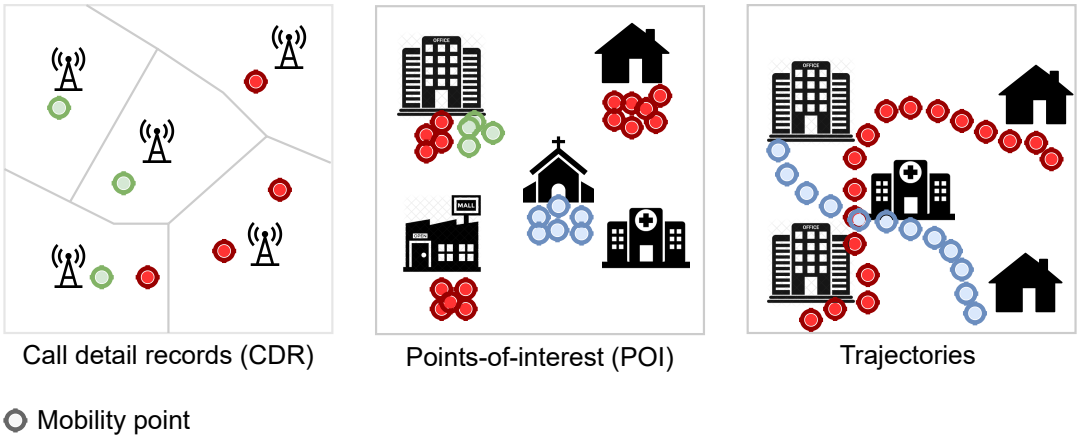


**Figure 2.1:** *Different representations of mobility data*

### Mobility Datasets

A mobility dataset is a set of mobility traces of multiple users. There are various sources to obtain mobility datasets, either public or private. Examples of public repositories giving access to mobility datasets include the Crawdad project [213] and the Safecast project [46]. Private mobility datasets are generally provided on-demand. They are collected by manufacturers of devices with a geolocation system, by LBS providers, geosocial network APIs (*e.g.*, Twitter), or by establishing agreements with telecom operators to get a call detail records dataset (CDR). All these datasets are pseudonymized, *i.e.*, the real identity of users is replaced by a random identifier.

In this thesis, for our experimental evaluation, we mainly used Geolife [257] that contains the mobility of users in the city of Beijing, MDC [135] that contains the mobility of users in the city of Lausanne, Privamov [169] that contains the mobility of users in the city of Lyon and Cabspotting [213] that contains the mobility of cab drivers in the city of San Francisco. In addition, we request two private datasets: the first one contains air pollution measurements in the city of Toulouse collected from bike trips [31] and the second one is a CDR dataset gathered by a major telecom operator in Shanghai, China [1].

### 2.1.2 Threats on Mobility Data

The mobility data faces multiple threats, which differ in their adversary knowledge and goals [238, 199]. The work of Wernke *et al.* [238] classifies threats on mobility data according to the attacker's knowledge: (1) Whether the attacker has access to a single or multiple mobility records. (2) Whether the attacker has access to only contextual information such as traffic statistics, road maps, and yellow pages. Each piece of information can help the attacker to target a specific attribute in the mobility record *(identity, location, time)*. For instance, the attacker can track the exact location of a user from multiple obfuscated locations or areas (location attacks [219]) by leveraging time to derive information such as the speed of the user's movement [98] or by analyzing the contextual information attached to her mobility [220].

Differently in the work of Primault *et al.* [199], the authors categorize the location privacy threats into four practical threats according to the perspectives and objectives of the attacker. In the following, we describe each one of them.

#### Point of Interest Inference

Points of interest (POIs) represent meaningful locations where a user stays for a significant time, such as home, workplace, or worship place. Many techniques are developed in the literature to extract POIs. Both heuristics and clustering algorithms are examples of such techniques [260, 123, 84, 131, 170, 216]. POIs are particularly sensitive as they can easily reveal valuable personal information about users, such as their hobbies, gender, religious beliefs, sexual orientations, political preferences, and health status. For instance, Figure 2.2 illustrates the mobility of a user in Paris. By just visually analyzing the spatial mobility of the user, we can infer the home location and deduce the religious beliefs of that user as he visited the big mosque of Paris. Indeed, more powerful tools exist to get the exact home address thanks to reverse geocoding [4] combined with temporal analysis to get accurate results.

In the same direction, Gambs *et al.* [85] showed how the home address of some taxi drivers in the city of San Francisco can be inferred. They validated their results thanks to satellite views which show where yellow cars are parked. Also in [79], the

authors can distinguish between Muslim and non-Muslim taxi drivers by correlating their movements with the five prayers of the day. The authors in [133] design Placer to classify locations semantically by giving significant labels (*e.g.*, home, workplace, shopping center, hospital, *etc.*,). The model reached up to 74% accuracy thanks to "home" and "workplace" which are the most visited locations. Keles *et al.* [125] use a bayesian model that considers the duration of stay at a POI, the day of the week, and the arrival time to predict the category of the POI.

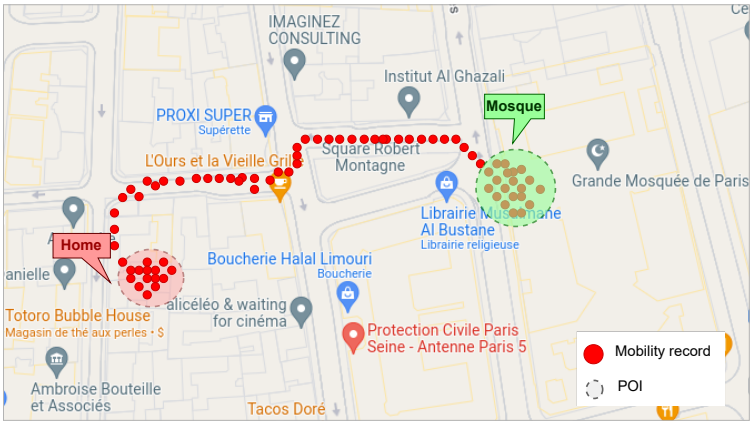


**Figure 2.2:** *Example of POIs inference.*

### Social Relationships Inference

Mobility data is also used to discover social ties and interactions between users. Generally, it states that if two users visit a location at similar moments, they are likely to be socially linked. For instance, family members stay at the same house at night, colleagues work together in the same area during the day, and friends meet at the same bars/ restaurants on weekends. Bilogrevic *et al.* [35] classify relationships between students of the EPFL campus using WiFi access points. They used two thresholds to detect stops and proximity of users and thus could detect meetings between users. Then, they built a classifier to infer the social links between students, whether they are classmates, friends, or others. They used a questionnaire and a database of courses as a ground truth. Similarly, Wang *et al.* [233] build a decision tree model where the interaction time and the physical proximity between individuals are used to classify social relationships between users. Moreover, many other works have been proposed to infer social relationships and demographic attributes relying on mobility

data combined with online or offline social networks [126, 188, 179, 145, 140, 141] to cite a few.

### Future Mobility Prediction

Predicting human mobility has received a lot of attention in the last decade. The research community has been actively proposing novel models to predict the mobility or activity of a user [251, 148]. These models generally rely on historical mobility data, transformed into diverse forms. For instance, Gambs *et al.* and many other works [86, 158, 54] have widely used Mobility Markov Chains (MMC) and their variants to model the movement habits of people and predict future human mobility. However, these models were later criticized for being ineffective over the long term. To handle this limitation, Gambs extended the work by considering $n$ past POIs instead of only a POI in the MMC. In addition, Sadilek and Krumm [209] propose far out, a system based on Fourier transformation and principal component analysis (PCA) to capture meaningful mobility patterns to predict long-term mobility (*i.e.*, months or years). Recently, with the advent of deep learning, Feng *et al.* [76] present DeepMove, a mobility prediction model based on recurrent neural networks (RNN) combined with a historical attention network to capture meaningful periodical patterns at different granularities. Also, Li *et al.* [137] design a long-short-term memory model (LSTM) with a hierarchical temporal attention model to predict future locations. For the same purpose, Dang *et al.* [61] propose a framework that builds a spatial-temporal embedding and a dual attentive network to learn sequential patterns in a trajectory and its correlation with other ones. Similarly, in [77], the authors present a predictive model based on spatial-temporal embedding and sequential modeling with LSTM in a federated learning approach.

### Re-identification Risk

The re-identification risk is defined as the ability of an attacker to associate an anonymous mobility trace to its originating user based on background knowledge previously built using past mobility data. According to the NIST, re-identification is the process of matching anonymized data and its originating user by leveraging publicly available information and auxiliary data. This thesis focuses on mobile user

re-identification risk as a means of location privacy risk assessment. Thus, we define and review the re-identification risk throughout the next section as a first step. Later on, we present related work on assessing the privacy risk using the re-identification threat in the context of mobility data in Section 2.3.1.

### 2.1.3 User Re-identification Risk

In the last decade, the research community has been remarkably active in mitigating privacy threats. While some researchers continuously enhance users' privacy by proposing powerful protection mechanisms to effectively anonymize sensitive data while preserving its utility [258], others keep discovering ways to break such anonymity [199]. A seminal work in the theme of re-identification is the work of Narayanan *et al.* [173] where more than 80% of users in the Netflix dataset (containing movies rating) are re-identified by matching them to a publicly available Internet Movie Database (IMDb). After that, the re-identification risk has affected several contexts where multiple kinds of data are at risk. For instance, some works study the re-identification risk on web search engines through individual search queries [92, 230], others on medical data [62, 202], social networks [174, 222, 117], or even on re-identifying programmers from their source code or executable binaries styloemetry [19, 47].

In the context of location privacy, mobility data has been the origin of many re-identification attacks. Several studies have been conducted on re-identifying users based on their mobility data. They generally move through two phases [152]: (i) a mobility profile construction phase and (ii) a re-identification phase, as depicted in Figure 2.3. In the first phase, the past mobility data of known users is transformed into meaningful representations, called mobility profiles. Then, once an anonymous mobility trace is received, the re-identification process consists of comparing its anonymous mobility profile to known ones (*i.e.*, previously built in the first phase) thanks to similarity metrics and retrieving the identity of the closest known profile. We formally define the re-identification risk in Equation 2.1. It is a function that takes as input an anonymous mobility trace $T$ and a set of historical mobility traces $\mathbb{H}$, known by the attacker (*i.e.*, background knowledge), and returns the identity of a user among a set of existing users $\mathbb{U}$.

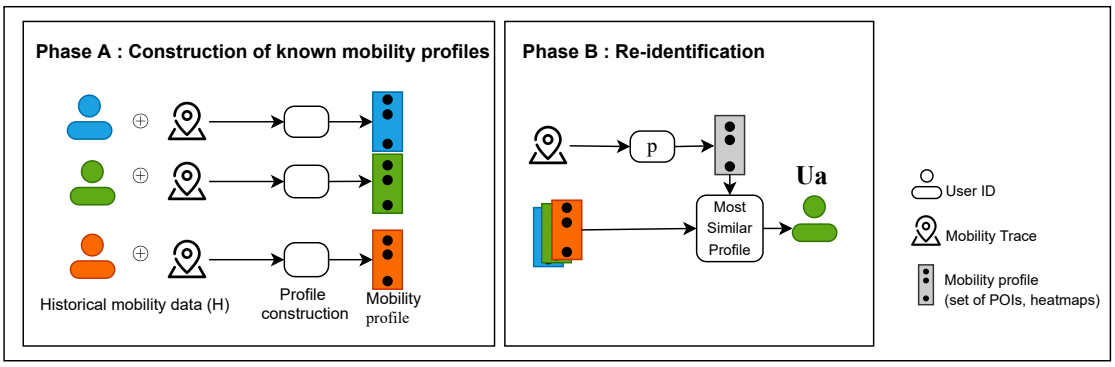


**Figure 2.3:** *Re-identification process.*

$$\begin{array}{rcccl} \mathcal{A} & : & (\mathbb{R}^2 \times \mathbb{R}_+)^* & \rightarrow & \mathbb{U} \\ & & T & \mapsto & \mathcal{A}(T, \mathbb{H}) = U_a \end{array} \tag{2.1}$$

In the literature, several mobility profiles are considered to extract discriminating mobility patterns from raw mobility data. Examples of mobility profiles are illustrated in Figure 2.4. Specifically, in the work of Gambs *et al.* [87], they model the mobility of users as Mobility Markov Chains (MMCs) where the states are POIs, and the edges represent the probability of transition between each pair of POIs. Multiple distance metrics are defined to compare MMCs. They are used to re-identify users by associating each MMC built from an anonymous trace to the identity of the closest MMC belonging to a known user. Moreover, Chen *et al.* [56] propose a density-based hidden Markov Chain model to capture hidden personalized mobility patterns by considering the spatial and temporal information of the mobility data. Primault *et al.* [194] characterize the mobility of a user as a set of POIs. The similarity between sets of POIs is based on the minimum geographical distance between their respective POIs. In the work of Naini [172], users are re-identified by matching their statistical characteristics between anonymous and known mobility data. They rely on a weighted bipartite matching algorithm to solve the problem of re-identification. In the same spirit of the previous work, Maouche *et al.* propose AP-Attack [152]. It is a re-identification attack that uses a heat map structure to represent the mobility traces of users. It divides a map into cells of approximately equal size. Each cell represents the proportion of time the user spent in that region. Then, they compare heat maps using the *Topsoe Divergence* metric [72]. In addition, the authors in [207] have developed a series of techniques for the re-identification of users in location-based

social networks. They mainly exploit the GPS coordinates of check-ins and the frequency of location visits to successfully recognize the user's identity. Moreover, Jin *et al.* formalize the re-identification problem as a k-nearest neighbor search based on similarity of mobility profiles [122]. Other works use side channels (*i.e.*, external information) to re-identify users. For instance, in the work of Srivatsa *et al.* [222], the authors use social network information such as contact graphs to re-identify users. They assume that a user may be recognized by those she meets. Also, in the work of Cecaj *et al.* [49], they present a probabilistic approach to re-identify users from anonymized CDR datasets by matching them to geo-referenced social networks data. Recently, the authors in [80] propose DART, a scalable framework to re-identify anonymous mobility traces by transforming them into sets of POIs using [101] and by associating these POIs to sparse known social trajectories thanks to spatio-temporal closeness scores. Furthermore, in the work of Massart *et al.* [157], the concept of re-identification attacks is revisited. They consider a combination of information-theoretic and security metrics to capture internal and external data leakage from a pseudonymized database to identify users.

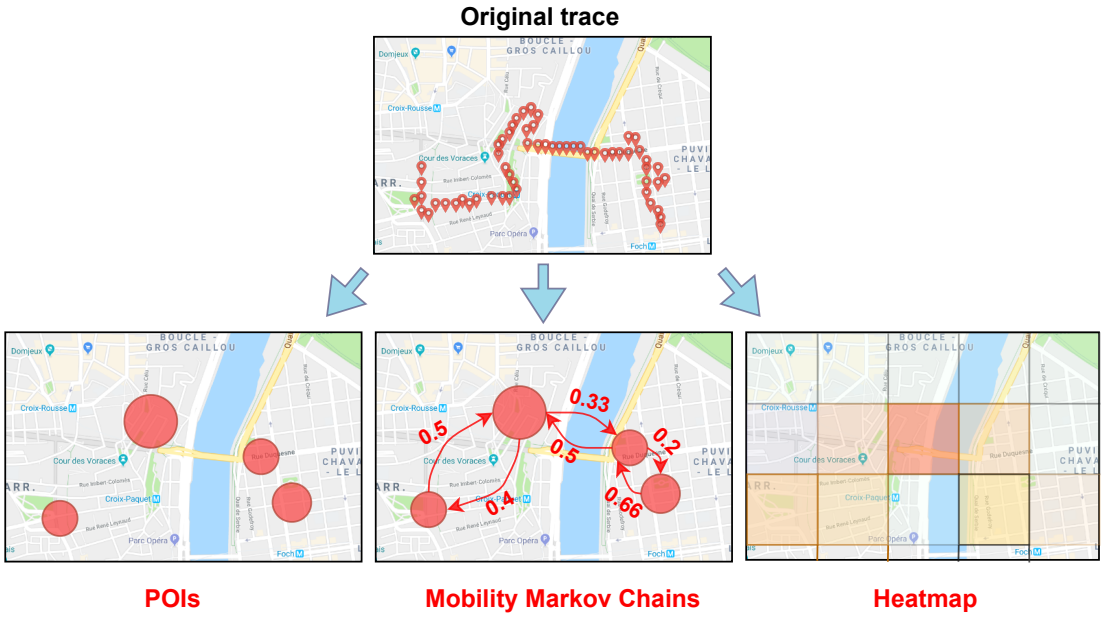


**Figure 2.4:** *Examples of mobility profiles.*

## 2.2 Related Work on Location Privacy Protection

Mitigating threats affecting location privacy has attracted the interest of many researchers. As a result, various location privacy protection mechanisms (LPPMs) have been proposed in the literature. Their main goal is to preserve user's privacy with a broad range of techniques such as perturbation, generalization, and fake data generation. More formally, an LPPM is defined as a function that takes as input one or multiple mobility records of $T$ and produces $T'$, an obfuscated version of this data.

We formulate an LPPM as in Equation 2.2, where $T \in (\mathbb{R}^2 \times \mathbb{R}_+)^*$ can be a single mobility record or a mobility trace, and $\Omega$ is a set of parameters.

$$\begin{array}{rccc} \mathcal{L} & : & (\mathbb{R}^2 \times \mathbb{R}_+)^* & \rightarrow & (\mathbb{R}^2 \times \mathbb{R}_+)^* \\ & & T & \mapsto & \mathcal{L}(T, \Omega) = T' \end{array} \tag{2.2}$$

### 2.2.1 LPPM Application Use Cases

There are various scenarios where LPPMs are applied to protect users' location privacy. According to Primault *et al.* [199], there are three types of LPPMs: online, offline, and semi-online, as illustrated in Figure 2.5. In the **online or interactive use case**, LPPMs obfuscate mobility data, record by record on the fly in real-time. The obfuscation process can be performed locally on the user's device or via a trusted proxy server (*i.e.*, anonymizer) or a P2P network. Thus, the LBS provider only receives the protected data and responds accordingly. Examples of applications that adopt online LPPMs include navigation, ride-sharing and POI retrieval systems, *etc.*,. The main challenge of these LPPMs is to operate rapidly so as to preserve the reactivity of the requested service. In the **offline or data publishing use case**, large amounts of mobility data, previously collected by the LBS provider or a telecom operator, need to be published for several purposes (*e.g.*, dataset analysis, dataset collection [213]). Therefore, the data publisher utilizes several LPPM techniques to ensure that the entire dataset is protected. These LPPMs can alter the user's data intrinsically, without the need for other users' mobility data, or exploit the knowledge

regarding other users to achieve some privacy guarantees such as k-anonymity. Finally, the **semi-online or crowd sensing use case** is an intermediate scenario where an LPPM is applied periodically on a batch of mobility records (*e.g.*, every hour) before reaching the service provider. It is possible that the user's response will be delayed until the anonymizer receives data from other users to run the LPPM effectively in this scenario. Once the service provider has received users' data, it performs different analyses and aggregations to inform users about their environment (*e.g.*, road traffic information, pollution maps, crowded places *etc.*,).

Furthermore, we have observed that some use cases are more constrained than others. Specifically, we can order them in the following way: offline → semi-online → online. Online LPPMs are more constrained than other types of LPPMs, simply because they are designed to handle real-time scenarios that require a responsive service. Using them for semi-online and data publishing use cases is possible. For instance, Huang *et al.* [118] propose a dummy-based mechanism that replaces the user's location with three randomly generated positions surrounding the real one. This mechanism is directly used on the user device for an online location searching use case and can be extended to offline or semi-online scenarios. However, an offline or semi-online LPPM can not be applied for an interactive scenario. For instance, Primault *et al.* [195] propose Promesse, a perturbation mechanism that uses a batch of mobility records to erase POIs. This can be used for an offline scenario. However, it is impossible for a online use case where the user only knows her current location.

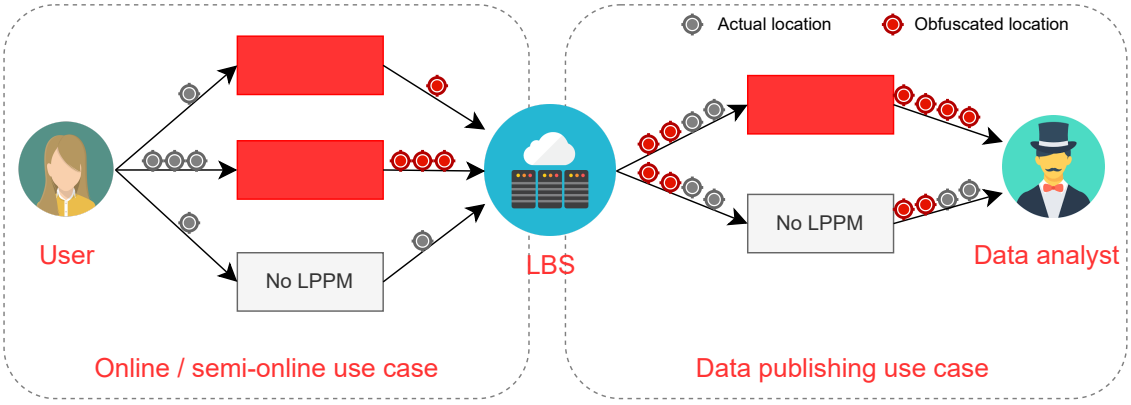


**Figure 2.5:** *The usage of LPPMs in different use cases*

### 2.2.2 LPPM Techniques

In the literature, there are a number of LPPMs based on various techniques to combat privacy threats affecting mobility data. Some of them need knowledge about the mobility of other users following a centralized or decentralized approach, while others do not need any external knowledge and are applied directly on the user's device. In this section, we review a non-exhaustive list of LPPM techniques.

#### Anonymization-based

Anonymization is the process of protecting the user's identity without altering their mobility data. To achieve this objective, pseudo-anonymization is first introduced by Gruteser *et al.* [105]. Essentially, it replaces the real identity of the user with a random and fictive identifier (ID), commonly referred to as a pseudonym. This process is managed by a centralized proxy server that knows all real IDs. As a result, it constitutes a single point of failure. In addition, this technique was criticized for being insufficient to ensure the privacy of the user's identity. Indeed, an attacker can easily link trajectories belonging to the same user or re-identify the owner of a mobility trace when the data is pseudonymized [75, 235, 64]. After that, Mix-zone is introduced by Beresford and Stajano [30] to improve pseudo-anonymization. A mix-zone is an area designed by an LPPM either statically or dynamically. In the mix-zone, the IDs of users are shuffled so that the LBS provider cannot track users' movements. In other words, when a user enters a mix-zone, the data is not shared and her ID is automatically mixed with the IDs of other users inside the same mix-zone. Several works are proposed to implement mix-zone mechanisms for online interactive use cases. Some of them focus on studying their optimal placement [81, 243], others rely on contextual information such as road network, and speed [181]. Additionally, Salas *et al.* [211] propose SwapMob as a mix-zone method for data publishing use case. In this work, the intersection of users' mobility traces is viewed as the location of mix-zones where IDs are exchanged. Thus, at the end of the process, a resulting mobility trace is made of multiple small segments of different users who have met during the day.

Unfortunately, the mix-zone mechanism has several drawbacks. Specifically, raw mobility data must be centrally processed via a trusted server. In addition, the

number of mix-zones placed in the city, has a considerable impact on the protection of users. In particular, if the number of mix-zones is low, an attacker may threaten the gathered segments when the user is outside a mix-zone. On the contrary, if the number of mix-zones is high, it may lead to a significant data loss (*i.e.*, little data is gathered since it is not shared when a user is inside a mix-zone). Furthermore, this approach is insufficient for user-centric analysis because mobility traces are shuffled between users. This may lead to compromised results.

#### Generalization-based

Generalization-based mechanisms have been widely used to ensure location privacy protection. In these techniques, instead of sending the exact location of a mobility record to the LBS provider, the LPPMs send coarse location information surrounding the user location aiming at enforcing k-anonymity, *i.e.*, making the user location indistinguishable among at least $k-1$ other users' location [224]. For this purpose, they either use cloaking regions in which the exact location of the user is hidden among at least $k-1$ different real users' locations, or they use dummy-based techniques in which fake locations are generated to simulate surrounding users.

*Cloaking Region.* It is first introduced by Gruteser and Grunwald [104] where the main idea is to report a coarse spatial region that contains k users instead of the exact user's location to the LBS provider. Numerous works on cloaking regions are proposed in the literature. For instance, Gedik *et al.* [91] propose Clique cloaking, a centralized solution based on spatial-temporal cloaking where a proxy server generates clique graphs of k users and can delay queries to achieve k-anonymity. Mokbel *et al.* [168] propose Casper, a centralized solution that allows users to specify their privacy requirements (level of k-anonymity) to create spatial cloaking regions using a tree-based data structure. The same authors propose a decentralized spatial cloaking where users collaborate in a peer-to-peer protocol to generate cloaked regions [58]. These methods are designed for cloaking single locations in an online interactive use case. That is why Abul *et al.* propose Never Walk Alone [14], and its extension W4M (Wait for Me) [15] to consider k-anonymity for trajectories. They guarantee that at each instant, there are at least k users walking inside a cylindrical volume of radius $\sigma$. Also, Gramaglia and Fiore [102] use a spatial distance to merge mobility traces with generalization zones to achieve k-anonymity. One interesting and different work to

cloak trajectories for data publishing purposes is to use semantic cloaking [25]. It is based on a machine learning approach to predict the most probable semantic label that will replace each location to release.

In summary, cloaking methods enable users to hide their precise location. However, they suffer from several vulnerabilities. For instance, a malicious entity may analyze the consecutive cloaked regions the user has visited and intersect them to track and re-identify the target user. In addition, these techniques need to be applied in areas with high density *i.e.*, high number of active users in the region. Otherwise, larger cloaking regions are created, which are useless for the data analyst. These techniques are based on trust of either a proxy server or neighbors in case of P2P networks.

*Dummy-based Mechanisms.* The concept of dummies is first introduced by Kido *et al.* [129]. It states that, instead of relying on users' actual location to achieve k-anonymity, multiple fake locations, called "dummies" are used to hide the user's actual location from the service provider. Many works are proposed in this direction. For instance, Shankar *et al.* present SybilQuery [217], a decentralized solution that generates fake trips suited for navigation applications. It creates $k-1$ fake trips using external knowledge (*e.g.*, databases of past traffic information, road maps, *etc.*,) while preserving the properties of the effective one (*e.g.*, length, semantics of starting and ending points). In the work of Bindschaedler *et al.* [36], they design a generative model of synthetic mobility traces that share statistical features with the real mobility traces used in the learning phase. Moreover, Huang *et al.* [119] propose trilateration (TRL), a new way to achieve k-anonymity by generating dummy locations locally on the user's device. Specifically, for each location $l$ the user actually occupies, TRL generates three random locations $l_1$, $l_2$, and $l_3$ that are within a range of $r$ from the actual one. After that, the fake locations are sent to the service provider instead of the real location. This method is commonly used in location-based search services where trilateration is used to calculate the exact distance between the user's actual location and the results related to the fake ones.

The main advantage of dummy-based mechanisms is that users can generate dummies without the need of any trusted proxy server. However, their main challenge is the ability to produce realistic indistinguishable dummy locations. In the work of Peddinti *et al.* [185], the authors identify the fake trips with 93% accuracy when $k=5$ (4 fake trips for every real one) and past mobility is known to the attacker, which is

used to train a machine learning classifier. In addition, these methods suffer from a consequent overhead as they increase the quantity of data to process. For example, SybilQuery multiplies the number of trips sent to the LBS by its parameter $k$. In TRL, the trajectory length is multiplied by three, which increases the loading and the processing time on the server side. Also, this type of mechanisms is inadequate for some specific analyses, such as counting the number of unique users in a given place. This may lead to incorrect results.

#### Perturbation-based

Perturbation or alteration methods are a kind of LPPMs that alter mobility data either spatially (*i.e.*, latitude and longitude) or temporally (*i.e.*, timestamp) before reaching the LBS provider. The alteration can be done locally on the user's device without a privacy proxy server. It consists of adding random noise to mobility data. The amount of added noise has an impact on the privacy *vs.* utility tradeoff. In particular, the higher the amount of noise, the better the level of privacy. However, the perturbed data deteriorates the quality of the offered service on one side and becomes useless to a data analyst on the other side. Many related works based on perturbation have been presented in the literature. Pingley *et al.* [189] propose a context-aware privacy protection system for online usage. It transforms mobility data from two dimensions space (location and time) to single dimension space (Hilbert space), adds noise to the transformed data, converts it back to two dimensions space, and sends it to the LBS provider. In the work of Micinski *et al.* [162], mobility data is perturbed by reducing the precision of GPS coordinates. It truncates the decimals from both the latitude and longitude values. Moreover, Andres *et al.* [22] adapt the concept of differential privacy [69] in the context of mobility data. They propose Geo-indistinguishability (GEOI), a mechanism that perturbs the spatial information of data by adding Laplacian spatial noise to each GPS coordinate. The amount of noise is calibrated by a privacy budget $\epsilon$ (the lower the $\epsilon$, the higher the privacy level). This perturbation is done on the user's device without the implication of any external proxy server. Furthermore, the same authors propose an extension of $\epsilon$-GEOI [53], where they calibrate the amount of noise thanks to contextual information about users' environment. This method is effective against location attacks. However, if an attacker exploits the correlation between successive locations, the privacy budget $\epsilon$ loses its power to a $n * \epsilon$ ($n$ being the number of

mobility points). In addition, Primault *et al.* [195] propose a temporal perturbation of mobility data. It erases user's POIs by smoothing the speed of the mobility trace. It ensures that two successive mobility points are equidistant (*i.e.*, constant speed), and thus it is difficult for a malicious entity to guess where a user stopped or what their POIs are. In the same spirit, Maouche *et al.* [153] propose HeatMap Confusion LPPM (HMC) to protect trajectories. It represents the mobility trace of each user as a heat map. Each user's heat map is altered in such a way that it looks similar to that of another user. As a final step, it converts the altered heat map into a mobility trace by leveraging mobility traces of multiple users. Moreover, the authors in [206] use machine learning techniques (*i.e.*, generative adversarial networks) to obfuscate mobility data. They build a generator network to produce noise to perturb mobility data, and a discriminator classifier to evaluate the re-identification risk of the perturbed data. The above works (*i.e.*, [153, 206]) provide a reasonable privacy *vs.* utility tradeoff. However, they require centralizing raw mobility traces in a proxy server which is not desirable.

#### Protocol-centric LPPMs

The protocol-centric LPPMs are mechanisms that focus on setting up proper protocols to preserve privacy by design in mobile applications and services. Most of the solutions found in the literature make use of either cryptographic techniques or secure multi-party computation principles. Specifically, Mascetti *et al.* [155] propose two protocols to discover nearby friends in social networks. Users share cryptographic keys instead of disclosing their current location. The work of Ghinita [97] presents private queries for neighbor searching in LBS relying on a practical and optimized implementation of private information retrieval (PIR) [57]. It retrieves data items from a database without revealing the retrieved data to the LBS. In other words, if a user sends a query asking for a specific place or searching for nearby friends, the LBS provider will prepare the answer without knowing the user's request. Guha *et al.* [108] present KOI, a protocol involving two non-colluding servers where each of them holds parts of the user request content. Specifically, the former knows about user identities and locations but it does not know the association between them. The latter knows the mapping between anonymized users and encrypted locations but nothing about actual identities or locations. The system employs a privacy-preserving protocol that enables its two components to match without knowing the mapping between users and their

locations. In addition, the authors in [17] use homomorphic encryption combined with secure multi-party computation to build a privacy-preserving ride-sharing system. Furthermore, the work of kulkarni *et al.* [134] proposes a hardware-based technique where they developed a private location-based service using intel SGX[1]. The service retrieves the nearest POIs to the user's location. This process is performed within an enclave where all the exchanged messages (*e.g.*, user request containing her location, response) are encrypted. The authors argue that this approach is a promising solution, as it maximizes privacy with a high data utility and a marginal overhead.

In general, protocol-based mechanisms offer robust privacy guarantees for specific use cases. However, the computational complexity, and the overhead generated by these protocols are their main limitations. In addition, the applicability and integration of these protocols with existing systems and infrastructures is not an easy task.

#### User-centric LPPMs

In the last decade, the research community has observed that one-size-fits-all LPPMs are insufficient to protect all users' mobility data. For instance, fixing a similar value of "k" in k-anonymity methods or the same noise amount in perturbation methods may lead to overprotecting some mobility data and under-protecting the rest. As a result, the overprotected data decreases its utility needlessly, and the under-protected one is still subject to aggressive attacks. Maouche *et al.* [152] conduct an experiment where they launch re-identification attacks on mobility datasets protected with individual LPPMs. The results show that despite the protection with existing LPPMs, the percentage of re-identified mobility data can reach up to 78%.

That is why LPPMs need to be more adaptive to each user. Existing and emerging works propose user-centric approaches where each user's mobility trace is protected according to its sensitivity, characteristics, and preferences in terms of privacy and data utility. In the work of Maouche *et al.* [152], they propose HybridLPPM, a user-centric LPPM that protects each user mobility trace against re-identification attacks using a set of state-of-the-art LPPMs (*e.g.*, Promesse [195], GEOI [22], HMC [153], *etc.*,) with a predefined order. The order is given according to the degree of data

[1]Intel SGX: https://software.intel.com/en-us/sgx

distortion caused by each LPPM after obfuscating data. Finally, the LPPM that degrades mobility data the least is selected. At the end of this process, all users are protected by different LPPMs (*i.e.*, not only by one). Moreover, SmartMask [139] is a system designed to automatically learn users' privacy preferences under different contexts (*e.g.*, location semantic, frequency of visits, duration of visits, time). Once the privacy level is determined, different techniques of LPPMs obfuscate mobility data. LP-guardian is also an example of a user-centric solution [74]. It is implemented on Android users' smartphones, where a decision tree is used to choose the adequate action to perform against different threats. SmartMask and LP-guardian require richer datasets and not only timestamped mobility traces. Another complementary approach is to play on the configuration of LPPMs to protect the users' mobility data. To this end, it is more efficient to adapt an LPPM configuration to each user's behavior rather than considering an LPPM with the same configuration for all users' mobility data. In this direction, some authors exploited optimization algorithms to find the appropriate configuration for a given LPPM that ensures the tradeoff between privacy and data utility objectives. For instance, ALP is a framework that enables an automatic configuration of the LPPM parameters using simulated annealing [198]. PULP is another framework, which automatically configures LPPMs until reaching users' objectives in terms of privacy and utility [51].

### 2.2.3 Evaluating the Effectiveness of LPPMs

The effectiveness of LPPMs can be measured in terms of privacy, utility and performance. In this section, we provide a list of commonly used metrics to assess the effectiveness of LPPMs. This list is not exhaustive and additional metrics can be found in [199, 232].

#### Privacy Metrics

It corresponds to the evaluation of the privacy protection level offered by LPPMs relying either on theoretical metrics based on formal guarantees (*e.g.*, k-anonymity or differential privacy) or more practical ones.

*Theoretical Metrics.* Theoretical metrics are based on formal guarantees to

measure the privacy level of mobility data. At present, there are two main concepts that ensure a level of privacy protection offered by an LPPM: k-anonymity and differential privacy with their respective variants. Specifically, k-anonymity is first introduced by Samarati and Sweeney [224] in the context of database systems. It states that a user is hidden among at least $k-1$ other users with similar properties in the database. In the context of location privacy, this translates to cloaking a given user's exact location in a geographical zone where there is at least $k-1$ other users' location. The larger the value of k, the higher the level of privacy, *i.e.*, the probability of identifying the query initiator is under $\frac{1}{k}$. After that, many variants of k-anonymity are developed. For instance, l-diversity [149] and location diversity in particular [244]. The latter states that in a given cloaking region, there should be at least $l$ distinct semantic locations to prevent homogeneity attacks [67].

In addition, differential privacy ensures that the result of an aggregate query over a database should not be influenced by the presence or absence of a single element in the database [69]. In the context of location privacy, Andres *et al.* [22] propose Geo-indistinguishability, an instance of differential privacy to guarantee a certain level of location privacy. Through this concept, LPPMs attempt to protect the presence or absence of individual locations. Hence, their main goal is not anymore to hide that a user is part of a database but to hide where they have been. Thus, it is possible to control the reported locations with a privacy parameter $\epsilon$. The lower the value of $\epsilon$, the higher the level of privacy protection.

*Practical Metrics.* Unlike the theoretical metrics that rely on formal guarantees, practical metrics are measurements that assess how valuable data is after its obfuscation. Precisely, they measure what can be disclosed or leaked after the application of an LPPM. For instance, the resilience to adversary attacks has been used to assess LPPMs. In this direction, Primault *et al.* [195] use re-identification attacks to evaluate the re-identifiability of users before and after obfuscation. Moreover, POIs retrieval [195] is also a possible way to assess mobility data. POIs are sensitive pieces of information, and their disclosure may compromise users' privacy. That is why an LPPM that minimizes the retrieval of POIs after obfuscation is promising. Other additional metrics are discussed in the work Wagner *et al.* [232].

### Utility Metrics

There are two categories of utility metrics proposed in the literature to measure the quality of the generated data by an LPPM [199]: (i) **Data-centric or quantitative metrics** that measure the distortion between the original and the obfuscated mobility data. Examples of such metrics include spatial distortion in the work of Primault *et al.* [196], spatiotemporal distortion (STD) in the work of Maouche *et al.* [153] where a spatial error is calculated under a temporal constraint, and finally the area coverage metric (AC) [198] which computes the overlap between the obfuscated and the original mobility trace using the F1-score. (ii) **Application-centric or qualitative metrics** which compare the result of a given application before and after applying an LPPM. Examples of such metrics include the range queries metric, a classical operation that counts the number of unique users who go through areas during a time window before and after obfuscation [196]. Also, the work of Riboni *et al.* [203] measures the quality of venue recommendation before and after applying noise. In the same direction, the work of Boukoros *et al.* [42] measures the level of radioactivity before and after applying a defense mechanism in the context of mobile crowdsourcing applications.

### Performance Metrics

The protection of mobility data requires a lot of resources. We can measure the consumption of such resources using several metrics depending on the use case scenario. For instance, the running time of an LPPM is a significant performance metric in the case of real-time and interactive usage of mobile applications. Communication overhead and energy consumption are also critical elements to consider, especially when the processed data is exchanged through the network or the LPPM is implemented locally on the user's mobile device. Finally, scalability is a crucial aspect to evaluate. This is because existing LPPMs are applied increasingly by numbers of users and thus should be able to handle a large volume of data load.

In the remainder of this thesis, we define the metrics for privacy, utility, and performance used in every evaluation in its corresponding section.

## 2.3 Related Work on Location Privacy Risk Assessment

Human mobility data is a source of sensitive and personal information, and its misuse may lead to serious privacy violations, revealing many insights about an individual's private life. That is why data protection authorities (*e.g.*, CNIL [63]), regulators (*e.g.*, EU GDPR in Article 35 regarding risk assessment[2]), Data Privacy Officers (DPO) and the research community (*e.g.*, OWASP [240]) have been actively proposing algorithms, methodologies, and tools for privacy risk assessment. In the context of location privacy, location privacy risk assessment (LPRA) has evolved from qualitative decisions [65] to quantitative metrics that allow better quantifying the privacy bounds of human mobility and thus improving the privacy policies for protecting that data. In the following, we present two practical means of LPRA, namely re-identification and uniqueness assessment.

### 2.3.1 Re-identification as a Means of Privacy Risk Assessment

In the context of location privacy, re-identification attacks have been widely used to quantify the privacy risk of re-identifying individuals. In the work of Pratesi *et al.* [192], they propose PRUDENCE, a system for assessing privacy risks in the data sharing ecosystems [192]. It empirically measures the probability of re-identifying an individual user in the dataset, using all possible background knowledge an adversary can collect. The system uses a perfect matching function between the data in each background knowledge and the tested user data and then returns the highest re-identification risk across all backgrounds. The proposed framework is effective, however, it has a high computational complexity as it considers all possible combinations of data to construct multiple backgrounds that the adversary might collect from the user mobility data. In addition to that, it centralizes raw users data. EXPERT is precisely proposed to overcome the aforementioned problem of computational complexity [175], through an extension of a previous system that allows privacy risk assessment of mobility data with a binary result (*i.e.*, low or

[2]The EU General Data Protection Regulation can be found at https://rb.gy/jntsfr

high) [186, 187]. For each user, EXPERT builds a mobility profile that captures individual mobility patterns extracted from a user's mobility data. Then, it uses a decision tree-based ensemble model to predict privacy risks. Despite the improvement made in EXPERT, it still requires centralizing raw data, as depicted in the left part of Figure 2.6.

### 2.3.2 Uniqueness as a Means of Privacy Risk Assessment

Another well-recognized metric to quantify privacy risks is the uniqueness assessment metric. It has been used in several fields such as web search [71, 247], fingerprinting for criminal investigation [248], and smartphone sensor fingerprinting [38]. In the context of mobility data, several techniques for privacy risk assessment have been proposed in the last years. A well-known study on the uniqueness of CDR data, collected from 1.5 million users, demonstrates that with only four mobility points, a user is highly distinguishable from other users [64] and can be re-identified in 95% of the cases. It means that human mobility acts as a fingerprint that uniquely identifies users, thus raising their awareness of privacy risks. An interesting study extensively analyzes anonymized CDR data, collected from 1.37 million users in 2000 applications, and demonstrates that the fingerprints of mobile application usage are highly unique, and most users are uniquely re-identified [228]. Other similar studies were conducted on GPS data [43, 208], on large-scale CDR data [250], on trajectories inferred from cyberspace cookie logs [234], or on POIs [48]. Furthermore, the uniqueness metric was also used for evaluating the effectiveness of anonymizing location data [221]. Moreover, uniqueness can also be related to k-anonymity [224] and precisely to historical k-anonymity where the historical data is used to construct anonymity sets, *i.e.*, a group of users sharing similar properties (*e.g.*, similar spatial and temporal information). In this context, the size of the anonymity set reflects how unique the user's data is. In this direction, Bettini *et al.* [32, 249] found that the history of mobility data can act as a quasi-identifier, which may uniquely identify a user and lead to serious privacy violations. Moreover, in the work of Xu *et al.* and Mascetti *et al.* [242, 154], instead of using the current locations of k neighbors of the query initiator in a location-based service, they exploit k historical locations of different mobile nodes previously collected to construct anonymity sets. In addition, Masoumzadeh *et al.* [156] consider a time window to achieve historical

k-anonymity. The main drawback of the above solutions is that they are based on a trusted proxy server (*i.e.*, anonymizer) that has a complete knowledge of all users mobility, as depicted in the left part of Figure 2.6. That is why decentralized architectures were proposed to achieve k-anonymity where the need of a trusted proxy server was not required anymore (central part of Figure 2.6). To this end, many related works rely on peer-to-peer communication between users to create anonymity sets [58, 95, 96]. However, these approaches generally assume that users are honest and share raw data between each other (e.g., their current location). More recent works use encryption techniques [94], secret sharing [78] or blockchains [245] to release the trust assumptions between users. However, the above decentralized solutions only exploit the current location of the user and hence can only assess instantaneous k-anonymity instead of historical k-anonymity.

In this thesis, one of our objectives is to estimate the uniqueness using historical k-anonymity without centralizing raw mobility data. We push our contributions to a federated learning architecture as depicted in the right part of Figure 2.6 and that we will present its background in the following section.

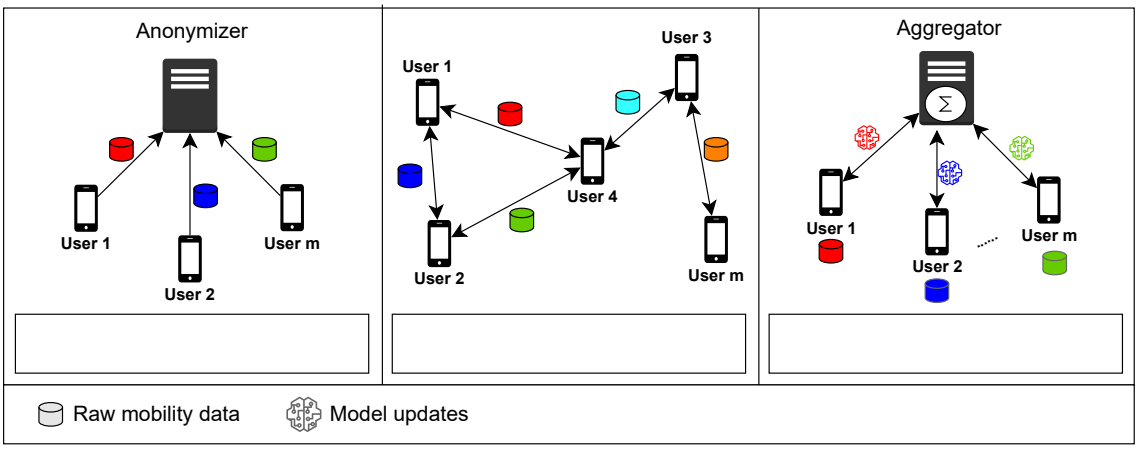


**Figure 2.6:** *Different architectures of uniqueness assessment systems*

## 2.4 Background on Federated Learning

In the following, we present a background on the federated learning paradigm (FL). We first define and describe the FL workflow in Section 2.4.1. Then, we present

its different architectures and types in Section 2.4.2 and Section 2.4.3, respectively. We finally discuss the main threats and attacks related to FL and list the common countermeasures to mitigate such threats in Section 2.4.4.

### 2.4.1 Definition and Workflow

Federated learning (FL for short) was initially introduced by [159]. It is a relatively new machine learning paradigm that enables training machine learning models on data from different sources without the need to store the data at a central server. FL is performed in several rounds where a set of clients (*i.e.*, known as workers) and a central server (known as the federator) are involved. At the beginning, as illustrated in Figure 2.7, the federator initiates the same model on all workers with either random parameters or predefined ones. For each FL round, the workers train locally the received model with their own local data to improve the machine learning model and send the updated model to the central server. The latter aggregates the received local models by averaging them, and produces a new version of the global model, which is sent back to the clients' devices. After that, a new FL round starts and the FL training process is repeated until the aggregate global model converges.

### 2.4.2 Federated Learning Architectures

We distinguish two main architectures of the FL protocol, namely: client-server architecture and peer-to-peer architecture. **Client-server architecture**, also known as the centralized FL, as illustrated previously in Figure 2.7. It requires a central server whose role is to initiate a global model, share it with clients for local training and wait for a predefined number of users (synchronously or asynchronously) to aggregate their models' updates to produce a new version of the global model. Nowadays, almost the implementations of FL follow client-server architecture, on top of them, the Google Gboard for Android [112, 55, 246]. The main advantage of this architecture is that it incurs low communication overhead. However, the central server presents a single point of failure and may leak sensitive information inferred from models' updates. We discuss these threats in Section 2.4.4. Unlike the centralized FL architecture, the **peer-to-peer architecture**, also known as the decentralized

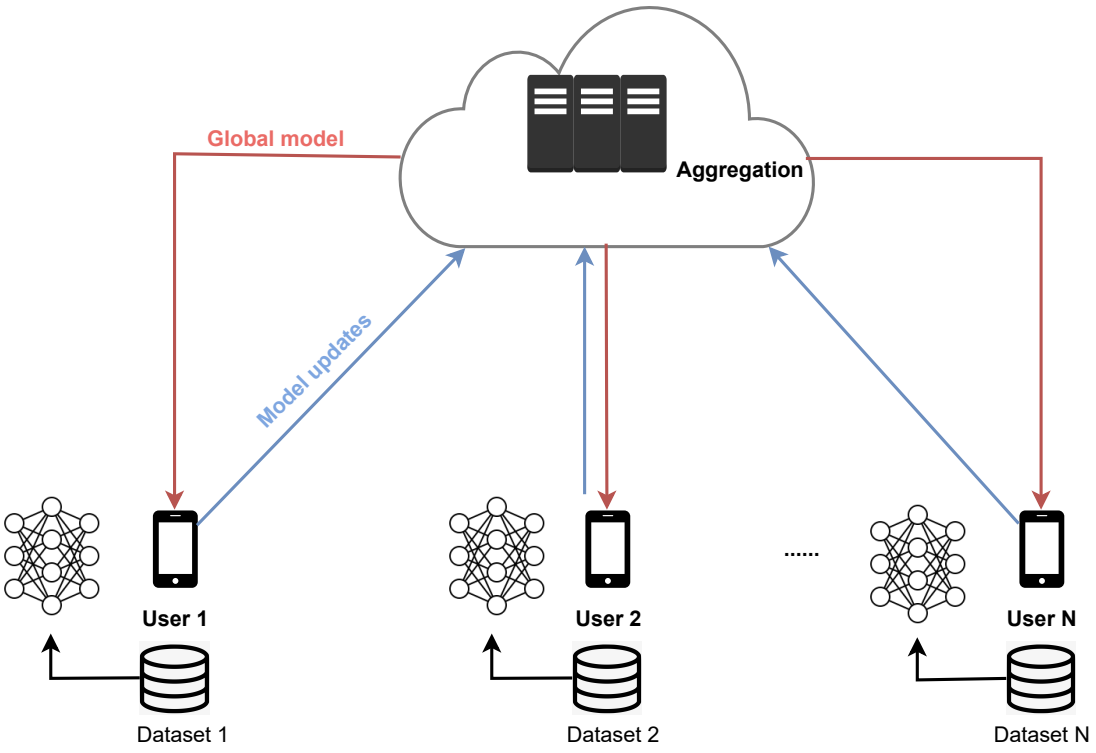


**Figure 2.7:** *A typical FL protocol*

FL, does not require a central server to maintain a global aggregate model. Instead, each client in the FL round has its own model and updates it thanks to information obtained from its neighbors [229, 147]. The downside of this architecture is that it leads to more communication overhead to achieve synchronization among multiple clients and opens the door to additional threats especially if the model falls between the hand of malicious parties. In this thesis, we opt for the *client-server architecture* to implement the FL protocol.

### 2.4.3 Federated Learning Types

There are three types of FL approaches; horizontal FL (HFL), vertical FL (VFL), and federated transfer learning (FTL). This classification is based on the data structure held by the clients, which may differ in the sample space and/or the feature space. Specifically, in the HFL, datasets owned by each client share the same feature space but differ in the sample space. For instance, the datasets from different hospitals

represent the same feature space, *i.e.*, the same attributes (*e.g.*, name, age, address, blood pressure), but of different samples, *i.e.*, each hospital holds data of its patients. In contrast, the VFL deals with clients whose data structure has the same sample space but with a different feature space. This category of FL is suitable to train a machine learning model for two or more organizations that have the same clients but hold distinct data types. Finally, the FTL involves datasets that differ in both the sample and feature space [144]. In this thesis, as users have different samples in the same feature space (*i.e.*, spatial or temporal features *etc.*,), we opt for the HFL to implement the FL protocol.

### 2.4.4 Threats and Countermeasures

In the last decade, the FL paradigm has attracted a widespread attention from both the research community and industry. It has significantly improved data privacy of multiple clients while benefiting from a global machine learning model in a privacy preserving way. Although the FL paradigm can avoid direct data leakage, it faces several vulnerabilities from both the central server and the clients. In the following, we provide a non-exhaustive list of attacks and threats affecting FL and present countermeasures to mitigate them.

#### Threats and Attacks in FL

There are mainly two categories of attacks which may occur during the training or inference phase of the FL protocol: (i) poisoning attacks and (ii) inference attacks. These attacks can be conducted by a curious/malicious user and/or the central server.

**Poisoning Attacks.** There are two categories of poisoning attacks: model poisoning and data poisoning. They are both performed by compromised clients. Specifically, in model poisoning attacks, the attacker tries to affect the FL model performance without being noticed during the local training process. The latter manipulates and modifies the model parameters either to minimize the accuracy of the global model on any test input [107, 73] or to misclassify specific inputs while maintaining high accuracy for the rest of the testing data [34, 23]. In data poisoning attacks, the attacker tries to affect the aggregated model fraudulently by injecting poisonous data

during the local data collection phase [215, 210, 218]. Compromised clients achieve this objective either by changing the data labels to predict an adversary's desired class (*e.g.*, labeling 20 km/h speed limit images as 80 km/h [218]) or by making imperceptible changes to the training data samples [210].

**Inference Attacks.** In the FL protocol, inference attacks are generally performed by a compromised central server. They aim to infer private information about users from their gradients. Indeed, a gradient may reveal sensitive information since it is derived from the client's local private data, For instance, it may reveal sensitive locations visited by a user in the case of human mobility prediction [77, 136]). There are various attacks proposed in the literature. Some of them try to determine if a specific training sample is used in the training process. This is known as membership inference attacks [201, 176, 143]. Others try to discover properties, specific attributes, or class representatives of the training samples in the FL protocol [115, 261, 88].

### Privacy Preserving FL

Privacy-preserving FL is an emerging research area. Many protection methods are developed to prevent the privacy issues affecting FL. These methods target either a compromised aggregator or compromised clients. Specifically, the compromised aggregator might be curious/malicious to infer sensitive and private information from users' updates, as presented in the previous section. To mitigate these threats, perturbation is an example of privacy-preserving methods. It adds noise to users' updates using differential privacy [69]. Many works are proposed in this direction [227, 214, 171]. However, their main drawback is their impact on the accuracy of the global model, which might deteriorate the quality of the proposed service. To avoid this issue, the secure aggregation primitive is used to mask users' updates so that the federator has only access to the aggregated model instead of individual model updates [255, 184, 124] and thus, the accuracy of the global model is not affected. In the same spirit, Bonawitz *et al.* show that participating users in FL can act as parties involved in the secure multi-party computation (SMC) [39]. Similarly, zhang *et al.* [254] apply SMC with homomorphic encryption to preserve the privacy of individual contributions, and the aggregator can only decrypt the global model once a threshold number of users have shared their updates. Although SMC provides good privacy guarantees in FL, its main drawback is its high communication

overhead. Furthermore, clients in the FL protocol might be curious/malicious for either inferring sensitive information or altering the global model to produce their desired result. To handle these users, the aggregator analyses users' updates by leveraging different techniques. Some of them are based on distance measurements, such as cosine metrics to detect unusual poisonous updates [83]. Others rely on clustering [218, 37] or anomaly detection [20]. Finally, hardware solutions are also a promising alternative, which consists of using trusted execution environments (TEEs) in an end-to-end manner both on the client side using solutions such as DarkneTZ [165] or GradSec [161] and on the server side using solutions such as Sear [256]. In these solutions, all the model parameters are encrypted and are manipulated in clear only inside hardware enclaves. TEEs have the advantage of preserving the accuracy of the trained model but they still require strong trust assumptions regarding the installation of the TEE and has limited resources (*i.e.*, limited trusted memory).

# Chapter 3

# MOOD: a Centralized Approach for Location Privacy Protection

## Contents

## 3.1 Motivation

In Chapter 1, we presented the privacy threats engendered by sharing exponentially growing amounts of information regarding users' locations. Furthermore, we described in Chapter 2 a state of the art of the privacy threats affecting mobility data with a focus on user re-identification attacks as a means of location privacy risk assessment. We also presented a state of the art on location privacy protection mechanisms (LPPMs) which try to tackle the privacy threats relying on a wide variety of techniques [110, 199].

To evaluate the effectiveness of these techniques, a variety of privacy risk assessment metrics are usually used and the resilience against re-identification attacks is one of them. The more an LPPM is able to protect against re-identification attacks, the better. However, when LPPMs are evaluated against re-identification attacks the focus is generally put on the protection of *the crowd*, *i.e.*, protecting the larger proportion of users possible, and little attention is given to users that remain unprotected. This minority of vulnerable users may have uncommon mobility behavior which makes them easily distinguishable and re-identifiable.

Considering a set of state-of-the-art attacks and LPPMs at the disposal of a data privacy officer (DPO) aiming at the protection of a given dataset, the question that the latter may ask is: *What should be done with mobility traces that are subject to re-identification despite the use of LPPMs?*. A straightforward, and safe solution that the DPO may adopt is to delete these vulnerable mobility traces from the protected dataset. However, this solution would engender a large data loss (42% in average and can reach up to 95%, as presented in Section 3.2).

In this chapter, we propose MOOD (MObility Data Privacy as Orphan Disease), a centralized user-centric multi-LPPM system which aims at protecting the mobility of *orphan* users, *i.e.*, users that are not protected against re-identification attacks while using individual LPPMs. The originality of MOOD is that it combines off-the-shelf LPPMs and applies a fine-grained protection. The LPPMs' combination is realized with the application of various LPPMs on the same trace in the form of function composition, while the fine-grained protection implies the application of various LPPMs on contiguous sub-traces. MOOD's choices are driven by the resilience to state-of-the-art re-identification attacks and the data utility metrics set

by the DPO.

We evaluate MOOD on four real-world mobility datasets and compare its performance to the application of individual and hybrid LPPMs [22, 118, 153, 152]. The results of our experiments show that MOOD is able to protect users' data in a range between 97.5% and 100% on the four datasets with an acceptable data utility and negligible data loss in comparison to the best competitor (HybridLPPM) which protects users' data in a range between 64% and 95% on the same datasets with lower data utility.

The work proposed in this chapter has been published and presented in Middleware Conference 2019 [127].

**Roadmap** The remainder of this chapter is structured as follows. First, we illustrate the handled problem with numbers in Section 3.2. Then, we present the design principles of MOOD in Section 3.3 and a detailed description of its component in Section 3.4. Further, in Section 3.5, we proceed to the experimental evaluation of our solution. Finally, we conclude this chapter in Section 3.6.

## 3.2 Problem Illustration

In this section, we want to showcase the ineluctable risk of re-identification of even protected mobility data with existing state-of-the-art LPPMs and measure the data loss caused by a conservative policy which relies on data suppression.

To this end, we consider a DPO that has to protect a given mobility dataset before its publication. The latter has access to a set of LPPMs and a set of user re-identification attacks found in the literature. In order to assess the effectiveness of the LPPMs in front of the attacks, the expert may decide to run the re-identification attacks on the protected dataset and choose the LPPM that better protects her original dataset. We performed such an experiment on four real-world mobility datasets protected using three state-of-the-art LPPMs (*i.e.*, Geo-I [22], TRL [118] and HMC [153]) and a hybrid solution proposed in [152] on which we ran three state-of-the-art attacks (*i.e.*, POI-Attack [193], PIT-Attack [87] and AP-Attack [152]).

The details of the used LPPMs and attacks are presented in Section 3.5. The results of this experiment are depicted in Figure 3.1. These results show, on each dataset, the number of users for whom at least one of the attacks was able to disclose their identities. From these results, overall the datasets, there are several users, from 19% to 88% that are not protected in front of re-identification attacks despite the use of single LPPMs.

The question that the DPO may ask in this situation is *what should be done with these vulnerable portions of the respective datasets?*. A safe answer would be to delete these parts of the datasets in order to prevent eventual user re-identifications that an attacker may perform on the published data. However, this may generate a massive data loss that ranges from 13% to 95% of the overall datasets, as depicted in Figure 3.2.

A closer look to the protected datasets shows that LPPMs perform differently from one user to another. Hence, a second step considered was to move to a user-centric approach where the hybridLPPM [152] is applied to each user of the considered datasets. The latter selects an LPPM among a set of LPPMs that resists to re-identification attacks (if any) with the best utility in terms of spatial and temporal distortion [153]. Column HybridLPPM of Figure 3.1 shows the ratio of non-protected users on the four datasets for which the best LPPM was chosen (*i.e.*, an LPPM that protects against all the three considered attacks with the lowest spatio-temporal distortion). This result shows that despite the use of an hybrid LPPM for protecting mobility datasets, there is still a large portion of users that are vulnerable to re-identification attacks (from 5% to 36%). Consequently, the generated data loss, as depicted in Figure 3.2 and that varies between 5% and 42% on the four datasets is still high.

The objective of this chapter is thus to design a novel methodology that combines off-the-shelf LPPMs to protect a given mobility dataset in front of a set of user re-identification attacks while minimizing the eventual data loss. In this way, we protect the *crowd* as it has been done in the literature and in addition, we provide other tools to protect *orphan* users.

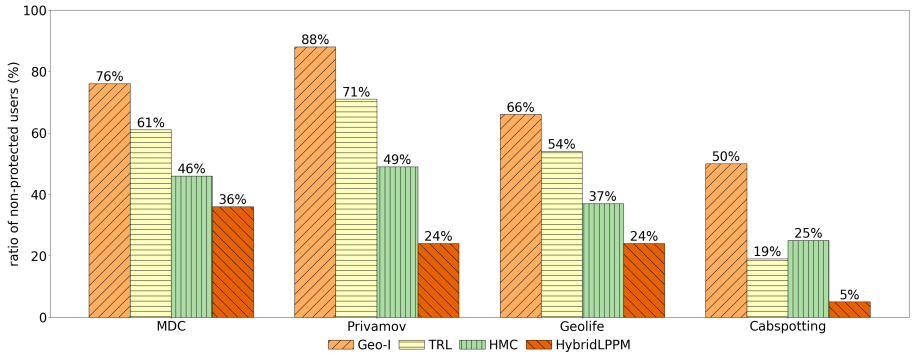


**Figure 3.1:** *Ratio of non-protected users with state-of-the-art LPPMs and Hybrid LPPM on four real-world mobility datasets*

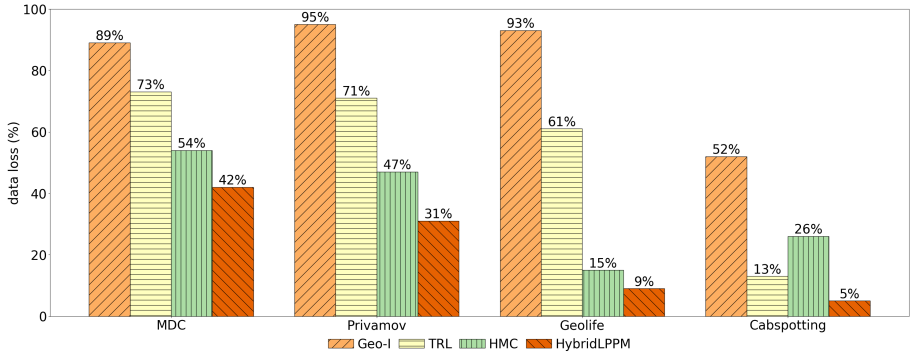


**Figure 3.2:** *Ratio of Data Loss with state-of-the-art LPPMs and Hybrid LPPM on four real-world mobility datasets*

## 3.3 MOOD Design Principles

In the following, we start by describing the system model of MOOD including the required background definitions in Section 3.3.1. Then, we present an overview of MOOD in Section 3.3.2.

### 3.3.1 System Model

Let $\mathbb{U} = \{U_1, U_2, ..., U_N\}$ be the set of users in the system. Each user is represented by two mobility traces, $T_{U_i}$ the one she wants to share and $H_{U_i}$ a past mobility trace used to control the risk of user re-identification. A mobility trace is a sequence of

mobility points including latitude, longitude and timestamp associated to a given user. To simplify, a mobility trace is considered as a time series $T \in (\mathbb{R}^2 \times \mathbb{R}_+)^*$.

**Definition of a User Re-identification Attack:**

To recall, the risk of user re-identification is formally defined in Equation 3.1 where $T$ is an anonymous mobility trace, $\mathbb{H}$ is a set of past mobility traces of known users and $\mathbb{U}$ is the set of users in the system.

$$\begin{array}{rcccc} \mathcal{A} & : & (\mathbb{R}^2 \times \mathbb{R}_+)^* & \rightarrow & \mathbb{U} \\ & & T & \mapsto & \mathcal{A}(T, \mathbb{H}) = U_a \end{array} \tag{3.1}$$

**Definition of a Composition of LPPMs:**

A composition of $p$ LPPMs $\{\mathcal{L}_{i_1}, \mathcal{L}_{i_2}, \ldots, \mathcal{L}_{i_p}\}$, *i.e.*, a subset of all available LPPMs in $\mathbb{L}$, noted $\mathcal{C}_p(\mathcal{L}_{i_k})$ is the application of $p$ LPPMs sequentially and gradually on a mobility trace. As described in Equation 3.2, it means that we start by applying the first LPPM $\mathcal{L}_{i_1}$. The resulting data is used as an entry for the second LPPM $\mathcal{L}_{i_2}$ and so on. The order of the LPPMs is important since it is similar to a composition of functions[1].

$$\begin{array}{rcl} \mathcal{C}_p(\mathcal{L}_{i_k})(T) & = & \mathcal{L}_{i_p} \circ \mathcal{L}_{i_{p-1}} \circ \ldots \circ \mathcal{L}_{i_2} \circ \mathcal{L}_{i_1}(T) \\ & = & \mathcal{L}_{i_p}(\mathcal{L}_{i_{p-1}}(...\mathcal{L}_{i_1}(T))) \end{array} \tag{3.2}$$

From $\mathbb{L}$, a set of $n$ LPPMs, the set of all possible compositions is denoted $\mathbb{C}$ where its size is given by the following expression:

$$|\mathbb{C}| = \sum_{i=1}^{n} \frac{n!}{(n-i)!} \tag{3.3}$$

[1] To simplify the notations we omit the parameters of each LPPM.

**Definition of a Fine-Grained Protection:**

The fine-grained protection splits the mobility trace into multiple sub-traces and protects each sub-trace independently with different LPPMs (from $\mathbb{L}$ or $\mathbb{C}$) as illustrated in figure 3.3. The objective of splitting traces is to separate discriminative mobility patterns. To this end, several techniques can be used, *e.g.*, splitting traces according to time, distance or inter-POIs.

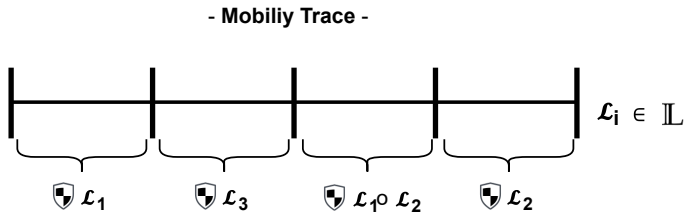


**Figure 3.3:** *Fine-grained protection of a mobility trace*

**Definition of an Orphan User:**

A user $U \in \mathbb{U}$ is an orphan user with respect to a set of LPPMs $\mathbb{L}$, a set of re-identification attacks $\mathbb{A}$ and a background knowledge $\mathbb{H}$, if she satisfies the property described in Equation 3.4, which states that for each one of the LPPMs, it exists at least one attack that successfully re-identifies the owner of the mobility trace $T_U$.

$$\forall \mathcal{L}_j \in \mathbb{L}, \exists \mathcal{A}_k \in \mathbb{A}, \mathcal{A}_k(\mathcal{L}_j(T_U), \mathbb{H}) = U \tag{3.4}$$

**Definition of a Protected User with Single-LPPM:**

A user $U$ is said to be protected by a single-LPPM if she satisfies the property in Equation 3.5, which states that there exists at least one LPPM in the set of considered LPPMs $\mathbb{L}$ that makes all the considered attacks in $\mathbb{A}$ fail at re-identifying the user.

$$\exists \mathcal{L}_j \in \mathbb{L}, \forall \mathcal{A}_k \in \mathbb{A}, \mathcal{A}_k(\mathcal{L}_j(T_U), \mathbb{H}) \neq U \tag{3.5}$$

**Definition of Protected User with Multi-LPPM:**

A user $U$ is said to be protected by multi-LPPM if she satisfies the property of the Equation 3.6. It states that it exists at least a composition of LPPMs that makes all the attacks fail at re-identifying the user where $\mathbb{H}$ is the background knowledge of the attacks.

$$\exists C_j \in \mathbb{C} - \mathbb{L}, \forall \mathcal{A}_k \in \mathbb{A}, \mathcal{A}_k(C_j(T_U), \mathbb{H}) \neq U \tag{3.6}$$

**Definition of Data Loss:**

We define the data loss over a dataset $\mathbb{D} = \{T_1, T_2, \ldots, T_N\}$, with the set of LPPMs $\mathbb{L}$ against the set of re-identification attacks $\mathbb{A}$ as the ratio of data size (counted by records) of non-protected mobility traces in $\mathbb{D}$. In other words, it is the amount of data which should be erased to avoid the re-identification risk. As described in Equation 3.7 (with $|\mathbb{D}|_r$ computes the number of records in $\mathbb{D}$).

$$\begin{aligned} &data\_loss(\mathbb{D}, \mathbb{L}, \mathbb{A}) = \frac{|\mathbb{D}_{NP}|_r}{|\mathbb{D}|_r} \\ &\mathbb{D}_{NP} = \{T \in \mathbb{D} \mid \forall \mathcal{L}_j \in \mathbb{C}, \exists \mathcal{A}_k \in \mathbb{A}, \mathcal{A}_k(\mathcal{L}_j(T_U), \mathbb{H}) = U\} \end{aligned} \tag{3.7}$$

### 3.3.2 MOOD Overview

MOOD (MObility Data Privacy as Orphan Disease) is a fine-grained multi-LPPM user-centric protection system. Its main objective is to protect the mobility trace of all users and in particular *orphan* users who are not protected by any available single LPPM. The architecture of MOOD is depicted in Figure 3.4 and its behavior is described in Algorithm 1. MOOD takes as inputs: the mobility trace of a user, denoted $T_U$, a set of LPPMs $\mathbb{L}$ of size $n$, a set of user re-identification attacks $\mathbb{A}$ of size $m$ and a utility metric $\mathcal{M}$. It returns obfuscated mobility data as an entire mobility trace $T'$ or as multiple sub-traces $\{T'_1, T'_2, \ldots\}$. It has three main components, the

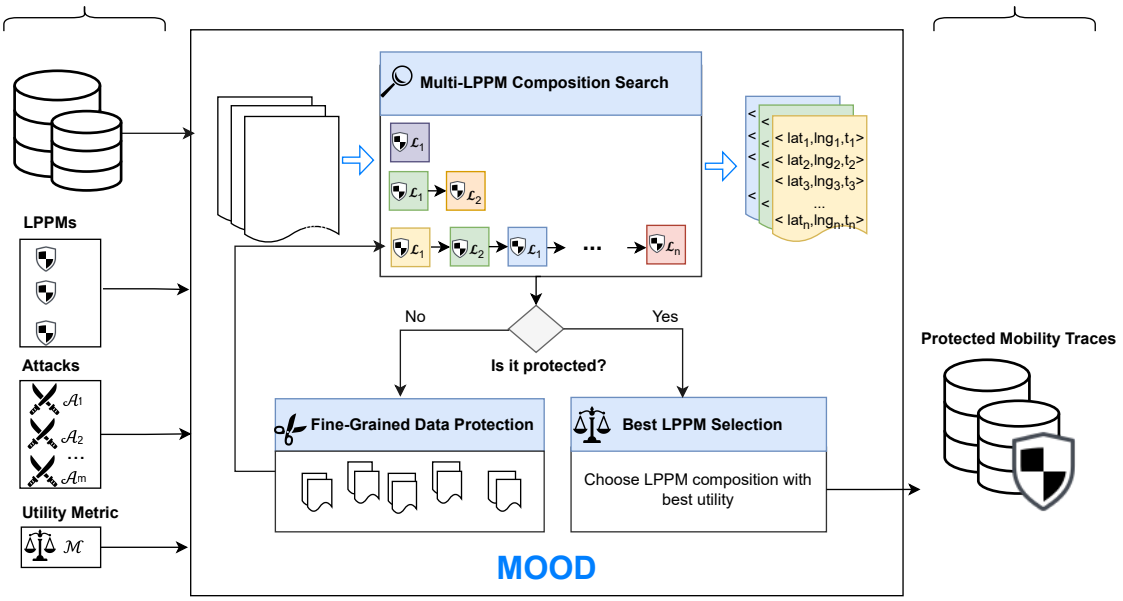


**Figure 3.4:** MOOD *architecture*

first component *Multi-LPPM Composition Search* aims at finding a composition of multiple LPPMs for orphan users. The second component *Fine-Grained Data Protection* manages mobility traces for which the first component was not able to find a protecting composition of LPPMs and uses the fine-grained protection. In this case, the latter splits the original trace into a set of sub-traces and sends each one back to the first component as depicted in Figure 3.4. Finally, in the last component *Best LPPM Selection*, only the protected mobility trace or sub-trace (*i.e.*, using either a single-LPPM or a multi-LPPM) against all the attacks with the best utility value is retained.

## 3.4 MOOD Detailed Description

In this section, we dive into a detailed description of each component involved in MOOD, namely, the multi-LPPM composition search (Section 3.4.1), the fine-grained data protection (Section 3.4.2) and the best LPPM selection (Section 3.4.3).

### 3.4.1 Multi-LPPM Composition Search

The Multi-LPPM Composition Search is the main component in our system. It takes as input the mobility trace of a user $T$, the set of LPPMs $\mathbb{L}$ and the set of re-identification attacks $\mathbb{A}$. First (lines 4 - 7), we start by applying LPPMs independently to search if there exists an LPPM that can protect the mobility trace against re-identification. Then (lines 15- 23), we apply all possible combinations of the considered LPPMs in an incremental and exhaustive manner so that the output mobility trace of the current LPPM becomes the input mobility trace of the next LPPM. For $n = |\mathbb{L}| = 3$, the number of different compositions is $|\mathbb{C}| = 15$ (given by Equation 3.3). After that, once a mobility trace $T$ is obfuscated by each composition of LPPMs, all the re-identification attacks $\{ \mathcal{A}_k \}_{k=1..m}$ are launched in order to evaluate the resilience of each composition of LPPMs and keep only ones that prevent from re-identification (if any). If all re-identification attacks fail in re-associating the obfuscated mobility trace $T'$ to its originating user, the privacy protection process is done and the user's mobility trace is protected by MOOD. In this case, the *Best LPPM Selection* component chooses the candidate that maximizes the utility metric (Section 3.4.3). However, if at least one re-identification attack succeeds, it means that the user is still vulnerable. In this case, the mobility trace of the user is undertaken by the following component.

### 3.4.2 Fine-Grained Data Protection

The Fine-Grained Data Protection is a complementary component in our system (line 28- 34). It is launched when the user's mobility trace is protected by neither a single LPPM nor a multi-LPPM. The idea we adopt is to split the original trace into a set of sub-traces and try to protect each sub-trace separately. For that purpose, several techniques can be used or combined for splitting the original trace, such as the fixed time slices where we split the mobility trace according to fixed times (*e.g.*, every hour) or to fixed distances (*e.g.*, every 1 km). In our work, we opt for the fixed time slices. The assumption behind going towards a fine-grained protection is that short mobility traces may contain less discriminating information than longer ones. Therefore, re-identification attacks which are based on profiling user mobility will be less successful at re-identifying users because the discriminating mobility patterns

**Algorithm 1** MOOD algorithm.

**function** $\mathcal{MOOD}(T_U,\mathbb{L},\mathbb{A},\mathbb{C},\mathcal{M},\delta)$
  $distortion \leftarrow \infty$
  $out \leftarrow \emptyset$
  **for** $\mathcal{L}_j$ **in** $\mathbb{L}$ **do** — Single-LPPM Protection
    $T' \quad \mathcal{L}_j(T_U)$
    $k \leftarrow 1$
    $limit \quad |\mathbb{A}|$
    **while** $\mathcal{A}_k(T') \neq U$ **and** $k \leqslant limit$ **do**
      $k \quad k+1$
    **end while**
    **if** $k > limit$ **then** $out \leftarrow out \cup \{T'\}$
  **end for**
  **if** $out \neq \emptyset$ **then**
    **return** $\{\arg\max_{T' \in out}(\mathcal{M}(T_U, T'))[0]\}$
  **else** — Composition of multi-LPPMs
    **for** $\mathcal{C}_j$ **in** $\mathbb{C} - \mathbb{L}$ **do**
      $T' \leftarrow \mathcal{C}_j(T)$
      $k \leftarrow 1$
      **while** $\mathcal{A}_k(T') \neq U \ \ \& \ \ k \leqslant limit$ **do**
        $k \leftarrow k+1$
      **end while**
      **if** $k > limit$ **then** $out \leftarrow out \cup \{T'\}$
    **end for**
  **end if**
  **if** $out \neq \emptyset$ **then**
    **return** $\{\arg\max_{T' \in out}(\mathcal{M}(T_U, T'))[0]\}$
  **else if** $length(T_U) \geq \delta$ **then**
    $S \quad Split\_in\_half(T_U)$
    — Fine-Grained protection
    $out \quad \emptyset$
    **for** $T_i$ **in** $S$ **do**
      $out \quad out \cup \mathcal{MOOD}(T_i, \mathbb{L}, \mathbb{A}, \mathbb{C}, \mathcal{M}, \delta)$
    **end for**
    **return** $renew\_Ids(out)$
  **else**
    **return** $\emptyset$
  **end if**
**end function**

collected from the mobility trace of the user are separated. This way, a user can still participate in the published dataset or in the crowd sensing campaign but only with multiple protected sub-traces that seem to come from different users. In practice, MOOD cuts the trace in half according to time and recursively calls for MOOD (line 32) with new user IDs in order to prevent a merging by the attacker (line 34). When the length of a mobility trace is shorter than $\delta$, the protection process for this trace is stopped and the corresponding mobility records are either erased from the published dataset or not sent to the crowd sensing server. The main role of the parameter $\delta$ is to stop the recursive split of traces. Moreover, in real use cases, the value of $\delta$ can be chosen according to the type of analysis the data will go through. For instance, in traffic congestion analysis (or count queries in general) there is no particular limit since the length of each sub-trace is not important to count the presence of users in particular places. But, if the application needs to study human mobility habits, a more reasonable value of $\delta$ is likely to be more than 24 hours.

### 3.4.3 Best LPPM Selection

It is important to protect data while maintaining high utility of the resulting trace. For that purpose, the Best LPPM Selection component is added to MOOD. Its main role is to choose $T'$, a protected version of a mobility trace with one among all the resilient LPPMs or multi-LPPMs against re-identification attacks, while maximizing the data utility. To this end, we rely on an utility metric $\mathcal{M}$ that measures the distortion of obfuscated data in comparison with the original data. The lower the distortion the better the quality of the resulting data. In MOOD, we measured the utility using the spatial-temporal distortion metric (STD) [153]. As defined in Equation 3.8, the STD is the average distance between each record of $T'$ and its temporal projection into $T$. The temporal projection of the record $x = (lat_x, lon_x, t_x)$ in $T'$ is the expected position $r_e$ in $T$ at time $t$. Specifically, we search for $r_i = (lat_i, lon_i, t_i)$ and $r_{i+1} = (lat_{i+1}, lon_{i+1}, t_{i+1})$ in $T$ such as $t_i \leq t_x \leq t_{i+1}$, then we compute $r_e$ the interpolation with the ratio $(t_x - t_i)/(t_{i+1} - t_i)$.

$$\mathcal{STD}(T, T') = \frac{1}{|T'|} \sum_{x \in T'} d_{temporal\_projection}(x, T) \tag{3.8}$$

## 3.5 Experimental Evaluation

In the following, we evaluate the effect of MOOD on the protection against re-identification attacks. In Section 3.5.1, we describe the experimental environment and configuration settings for the LPPMs and attacks considered in the experiments. Then, we describe the datasets in Section 3.5.2. To better understand the impact of the composition of LPPMs and the fine-grained protection on *orphan* users, we evaluate these parts of our system separately. Specifically, we evaluate MOOD's composition effect against single and multiple attacks in Sections 3.5.3 and 3.5.4, respectively. Then, for the remaining unprotected users, we analyze the effect of MOOD's fine-grained protection in Section 3.5.5. We finally study the impact of MOOD in terms of data utility and data loss in Section 3.5.6.

To resume, our evaluation answers the following questions:

- What is the effect of MOOD's multi-LPPM composition search in comparison to competitors on the protection against one re-identification attack? (Section 3.5.3)
- What is the effect of MOOD's multi-LPPM composition search in comparison to competitors on the protection against multiple re-identification attacks? (Section 3.5.4)
- How does the fine-grained protection handle orphan users? (Section 3.5.5)
- What is the impact of MOOD on data utility and data loss? (Section 3.5.6)

### 3.5.1 Experimental Setup

All the experiments were carried out in a computer running an Ubuntu 16.04 LTS OS with 5GB of RAM and 4 cores of 1.8Ghz each. The chosen LPPMs and attacks to conduct the experiments were taken from an open-source library [200] or the authors' own source code.

### User Re-identification Attack Configuration

In our experiments, we have considered three re-identification attacks, namely: AP-attack, POI-attack, and PIT-attack. Each attack has a set of parameters, described below. POI-Attack and PIT-Attack have two parameters for the extraction of POIs from mobility traces [259]. These parameters are the diameter of the clustering area and the minimum time spent inside it. These parameters are respectively set to 200 meters and 1 hour to accommodate small traces. AP-Attack has a configuration parameter that corresponds to the square cell size. It was set to 800 meters (default value in [152]).

In the evaluation scenario of the multi-LPPM composition search, we first considered AP-attack as the most aggressive attack among the chosen ones. Then for the rest of experiments, we combined all the attacks to form their union. We suppose that the DPO knows the ground truth about users' identities as the latter has access to all users traces stored in a central entity, known as the privacy proxy server (*i.e.*, anonymizer).

### LPPM Configuration

To evaluate MOOD, we selected three representative LPPMs, namely: (1) Geo-indistinguishability (GEOI) [22], (2) Trilateration (TRL) [118] and (3) Heat Map Confusion (HMC) [153]. Each LPPM belongs to a class of protection methods. GEOI is a data perturbation-based mechanism, TRL is a dummy-based mechanism for online services, and finally HMC is a combination of perturbation-based and dummy-based techniques.

Each LPPM has its own configuration parameters. These parameters have an impact on the privacy *vs.* utility tradeoff. In our experiments, we chose medium values of parameters because the objective of our study is not to find the best configuration as previous works did [50, 197] but to show that it is possible with a reasonable configuration and a relevant combination of LPPMs to reach an adequate tradeoff between privacy and utility. Specifically, GEOI has $\epsilon$ as a privacy parameter, which tunes the amount of noise added to the mobility data, (the lower the value of $\epsilon$ the higher the protection). We have fixed it to 0.01 which corresponds to a

medium privacy level. TRL has a radius $r$ from the real user's position where the fake locations are generated. The latter is set to 1 km. Finally, as HMC is based on heat maps, the cell size of the heat map is set to 800 meters as in the original paper [153]. Moreover, we compared MOOD to the HybridLPPM [152] with slight variations. Briefly, we selected the above LPPMs with the same configuration. Then, we ordered them according to the degree of data distortion they generate after obfuscation: (HMC $\rightarrow$ GEOI $\rightarrow$ TRL). Finally, we opt for the LPPM which degrades the least the mobility data while protecting it against re-identification, using the defined order.

### 3.5.2 Mobility Datasets

In our experiments, we used four real-world mobility datasets with a summary depicted on Table 3.1. These datasets are MDC [135], Privamov [169], Geolife [257], and Cabspotting [190].

In our experiments and for a fair comparison, we considered the 30 most active successive days of each dataset. After that, we split the mobility trace of each user chronologically into a period of 15 days used as a training dataset (*i.e.*, background knowledge) and the remaining 15 days used as a testing dataset (*i.e.*, data to be published). Only active users during those periods were considered.

Moreover, a mobility trace that is not protected by a Multi-LPPM in MOOD is split into sub-traces of 24 hours length before applying the recursive splitting algorithm presented previously. We choose chunks of 24 hours to simulate the scenario of a crowd sensing application where users send their data daily. Besides, we set the value of $\delta$ to 4 hours in MOOD's algorithm.

**Table 3.1:** *Description of datasets*

| Name | Cabspotting | Geolife | MDC | Privamov |
|---|---|---|---|---|
| # users | 531 | 41 | 141 | 41 |
| location | San Francisco | Beijing | Lausanne | Lyon |
| # records | 11 179 014 | 1 468 989 | 904 282 | 948 965 |

### 3.5.3 Evaluation of Resilience to a Single Re-identification Attack

As a first step in our evaluation, we want to showcase the problem of orphan users when a single attack is used by the DPO. We consider a set of state-of-the-art LPPMs ($n = 3$) and we select AP-attack as - the most powerful re-identification attack currently known in the literature - in order to evaluate the robustness of the generated data. We compared the result of MOOD's multi-LPPM composition search with the existing LPPMs applied individually on the four datasets. The results are depicted in Figure 3.5.

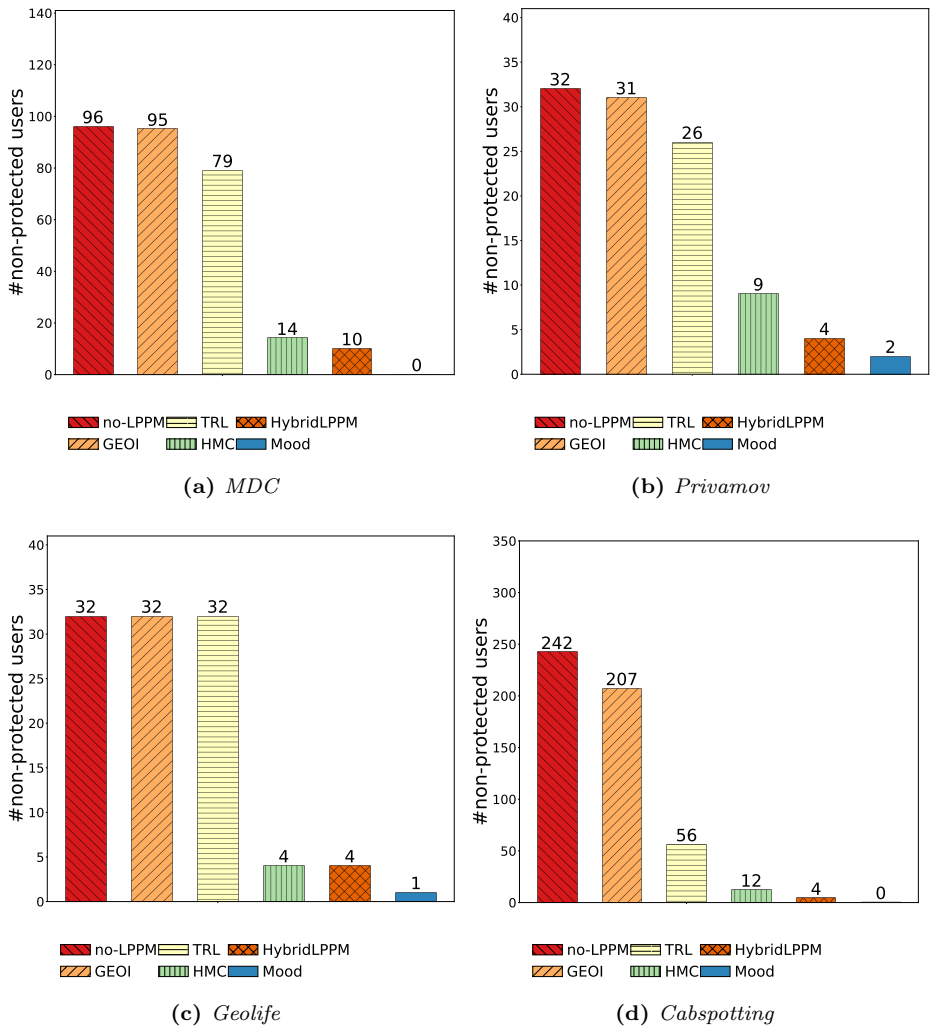


**(a)** *MDC* **(b)** *Privamov*

**(c)** *Geolife* **(d)** *Cabspotting*

**Figure 3.5:** *Resilience to one attack* – MOOD *vs. competitors*

In the MDC dataset (Figure 3.5a), 96 out of 141 users are re-identified when no LPPM is applied, which means that 45 users are naturally insensitive to AP-attack. Additionally, 95, 79 and 14 users are re-identified while applying GEOI, TRL, and HMC respectively. Whereas these numbers are lower (*i.e.*, 10 users) when HybridLPPM is used. All of those users are protected when using MOOD.

In the Privamov dataset (Figure 3.5b), 32 out of 41 users are exposed to re-identification threat. 31, 26 and 9 users are re-identified when GEOI, TRL and HMC are applied as a single LPPM respectively. Whereas in the case of HybridLPPM, only 4 users remain unprotected.

Similarly, in the Geolife dataset (Figure 3.5c), 32 out of 41 users are re-identified when no LPPM is applied against AP-Attack. Then 4 users are still unprotected by neither a single LPPM nor a HybridLPPM, whereas with MOOD's composition of LPPMs, only one user is still re-identifiable with AP-attack.

Finally, in the Cabspotting dataset (Figure 3.5d), nearly half of the dataset is naturally protected against AP-Attack (242 out of 536), this is due to the homogeneity of cab drivers moving patterns. After applying a single LPPM, almost all the remaining unprotected users became protected except 4 users, for which the application of MOOD with its multi-LPPM composition succeeds in protecting all of them.

### 3.5.4 Evaluation of Resilience to Multiple Re-identification Attacks

In this experiment, we consider a stronger virtual attacker where multiple re-identification attacks are used (*i.e.*, $m = 3$) to assess whether the protected users are uncovered by at least one of the attacks. This is possible because MOOD knows the ground truth about the real identity of the users. Indeed, all users' mobility data is stored in a privacy proxy server where MOOD is applied. The results are shown in Figure 3.6.

Specifically, in the MDC dataset, as depicted in Figure 3.6a, 107 out of 141 users are re-identified when no LPPM is applied. This means that 34 users are naturally

protected without the application of LPPMs. Thereafter, 107, 86 and 65 users are non-protected against at least one-attack among the considered ones when GEOI, TRL and HMC are applied individually. Then, 51 out of 141 users are still re-identified when HybridLPPM is applied. Whereas only 3 users remain non-protected with the MOOD's multi-LPPM composition.

In the Privamov dataset, as depicted in Figure 3.6b, 37 out of 41 users are vulnerable to the re-identification risk when no LPPM is used to protect the mobility data. Then, 36 users are re-identified when GEOI is applied. The latter is not resilient to the re-identification risk and its only way to be effective is to increase its level of privacy (*i.e.*, reduce its privacy parameter $\epsilon$) at the expense of data utility. Moreover, 29 and 20 users are non-protected when TRL and HMC are applied, respectively. These numbers decrease to 10 non-protected users when the HybridLPPM is considered. Finally, only 3 users remain re-identified with MOOD.

Similarly, in the Geolife dataset, as illustrated in Figure 3.6c, 32 users out of 41 users are unprotected against at least one among all the attacks. Then, the number of re-identified users decreases slightly to 28, 23 and 15 users when GEOI, TRL and HMC are individually applied, respectively. Furthermore, the application of HybridLPPM generated 10 non-protected users and finally, only 2 users are still vulnerable to the re-identification risk when MOOD is considered.

Finally, in the Cabspotting dataset, as depicted in Figure 3.6d, more than half of the whole users are re-identified in case of no LPPM. After that, 263, 131 and 65 users are re-identified when GEOI, HMC, and TRL are individually applied, respectively. Then, the number of re-identified users declines to 27 users with HybridLPPM. Lastly, while considering the multi-LPPM composition of MOOD, no users left unprotected. It means that MOOD is able to protect the whole mobility dataset including *orphan* users against all the considered attacks.

### 3.5.5 Evaluation of Fine-Grained Data Protection

As there are still few users who are vulnerable to the re-identification risk, we zoom on this category of users and apply the fine-grained data protection of MOOD. We start by splitting their mobility traces into multiple sub-traces of 24 hours length

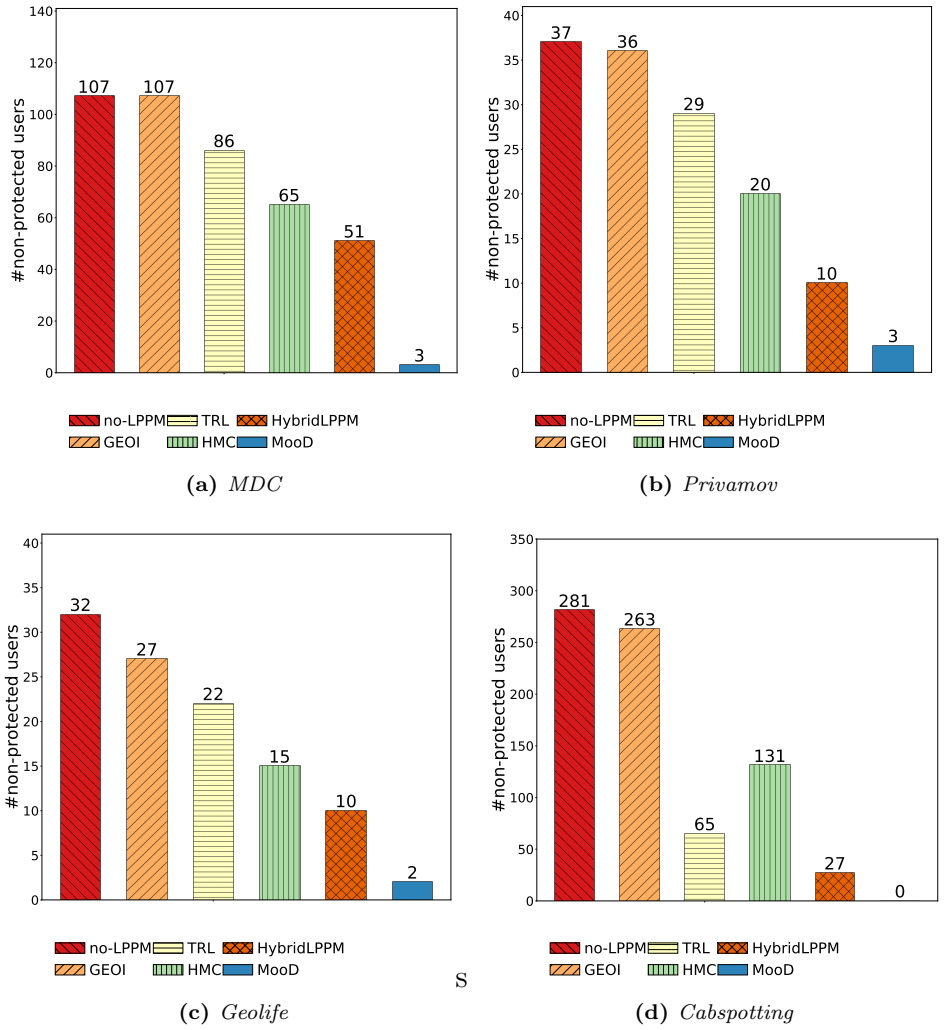


**(a)** *MDC* **(b)** *Privamov*

**(c)** *Geolife* **(d)** *Cabspotting*

**Figure 3.6:** *Resilience to multiple attacks* – MOOD *vs. competitors*

(as explained in Section 3.4.2). Then, each sub-trace feeds MOOD's multi-LPPM composition search component in order to protect it independently.

The results of Figure 3.7 illustrate how each user sub-traces are protected. Specifically, in the MDC dataset (Figure 3.7a), three users are not protected with MOOD's multi-LPPM composition, denoted $\{A, B, C\}$. Overall the three users, there are 68% of protected sub-traces with MOOD and the remaining sub-traces are still unprotected. Precisely, we can see that user $A$ became protected. User $B$ almost protected (92% of her sub-traces are protected), whereas User $C$ is still unprotected (only 11% of her sub-traces are protected). Thus, the granularity of the considered

traces has an impact on privacy protection.

With the Privamov dataset (Figure 3.7b), three users $\{D, E, F\}$ are not protected when the MOOD's multi-LPPM composition is used. After the fine-grained protection, the results show that 67%, 43% and 50% of the sub-traces of user $D$, $E$ and $F$ are protected respectively. Thus, the remaining users are partially protected. Finally, in the Geolife dataset (Figure 3.7c), only two users $\{G, H\}$ are not protected by MOOD's multi-LPPM composition. Then, after splitting their mobility traces, we obtain 4 sub-traces (*i.e.*, 2 sub-traces for each user), the results show that only one sub-trace is protected by MOOD.

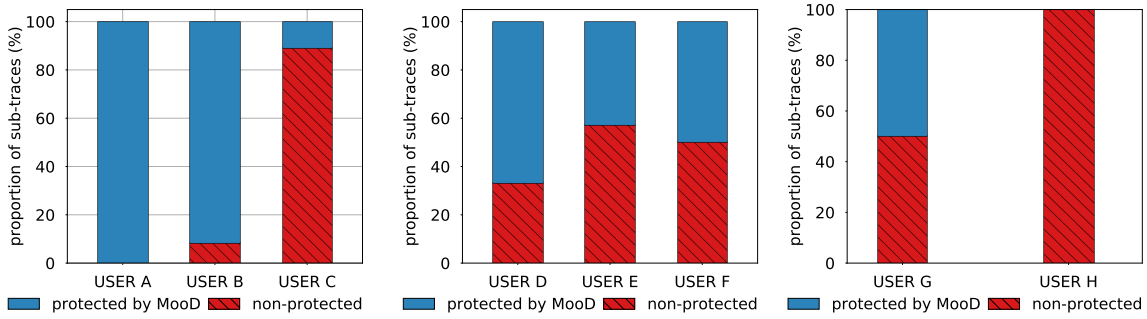


**(a)** *MDC* **(b)** *Privamov* **(c)** *Geolife*

**Figure 3.7:** *Fine-grained data protection with* MOOD

### 3.5.6 Evaluation of Mobility Data Utility and Data Loss

It is important to evaluate the effectiveness of MOOD in terms of data utility and data loss. In this chapter, as discussed in Section 3.4.3, data utility is measured using the spatio-temporal distortion metric [153] and the data loss is computed as defined in Section 3.3.1. After that, we compare MOOD to state-of-the art LPPMs, previously described in Section 3.5.1. We consider four different levels of the spatial-temporal distortion: low (*i.e.*, $<$ 500m), medium (*i.e.*, $<$ 1000m), high (*i.e.*, $<$ 5000m) and extremely high (*i.e.*, $\geqslant$ 5000m).

Overall datasets, as depicted in Figure 3.8, the results show that for all the protected users (*i.e.*, 754 users in the four datasets), 53.47% have a high utility using MOOD (*i.e.*, $<$ 500m) compared to its competitors, *i.e.*, 38%, 12%, 45% and 49%

with GEOI, TRL, HMC and HybridLPPM, respectively. Moreover, with medium utility (*i.e.*, < 1000m), MOOD outperforms its competitors with a ratio of 78% as to GEOI, TRL, HMC and HybridLPPM with 38%, 70%, 48% and 74%. This means that MOOD can provide a good balance between privacy and data utility in comparison to its competitors.

Depending on the degree of distortion, we can imagine several scenarios of data publishing and crowd sensing applications using MOOD. For instance, measuring the level of noise in a city when the distortion is low [150]. For medium distortion, MOOD can be used in an application that measures the level of pollution in a specific area. Finally, for high distortion, an application could be related to weather forecasting where the spatial precision of the protected data is not as sensitive as in the previous scenarios.

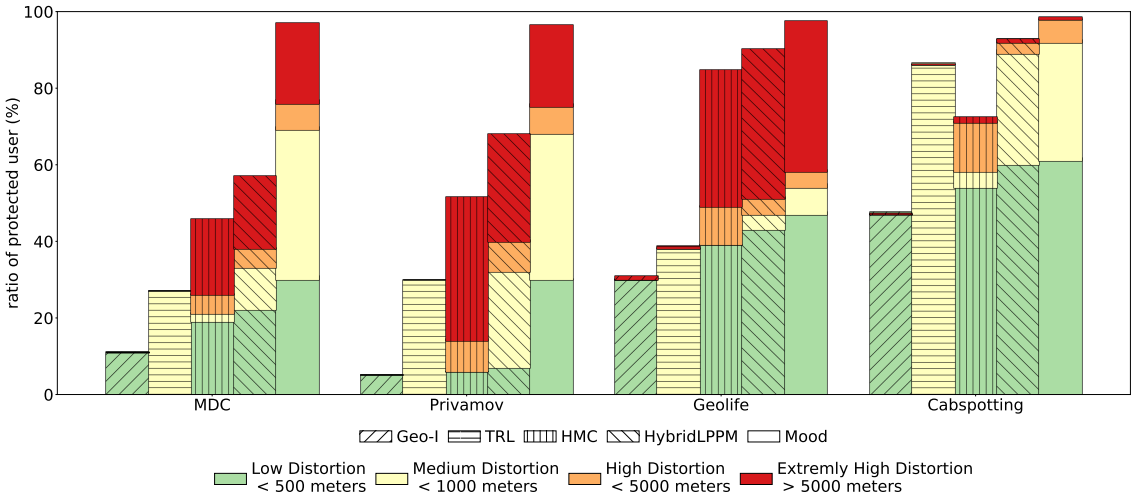


**Figure 3.8:** *Utility of protected data with* MOOD *vs. competitors*

In addition, we compared the data loss generated by MOOD and its competitors. The results are depited in Figure 3.9. In this figure, we found that a data loss between 14% and 95% is caused by the application of a single LPPM (*i.e.*, GEOI, TRL, and HMC). Furthermore, when HybridLPPM is used the generated data loss is between 5% and 42%. In contrast, MOOD generates a data loss between 0% and 2.5% which is a negligible amount of data compared to its competitors.

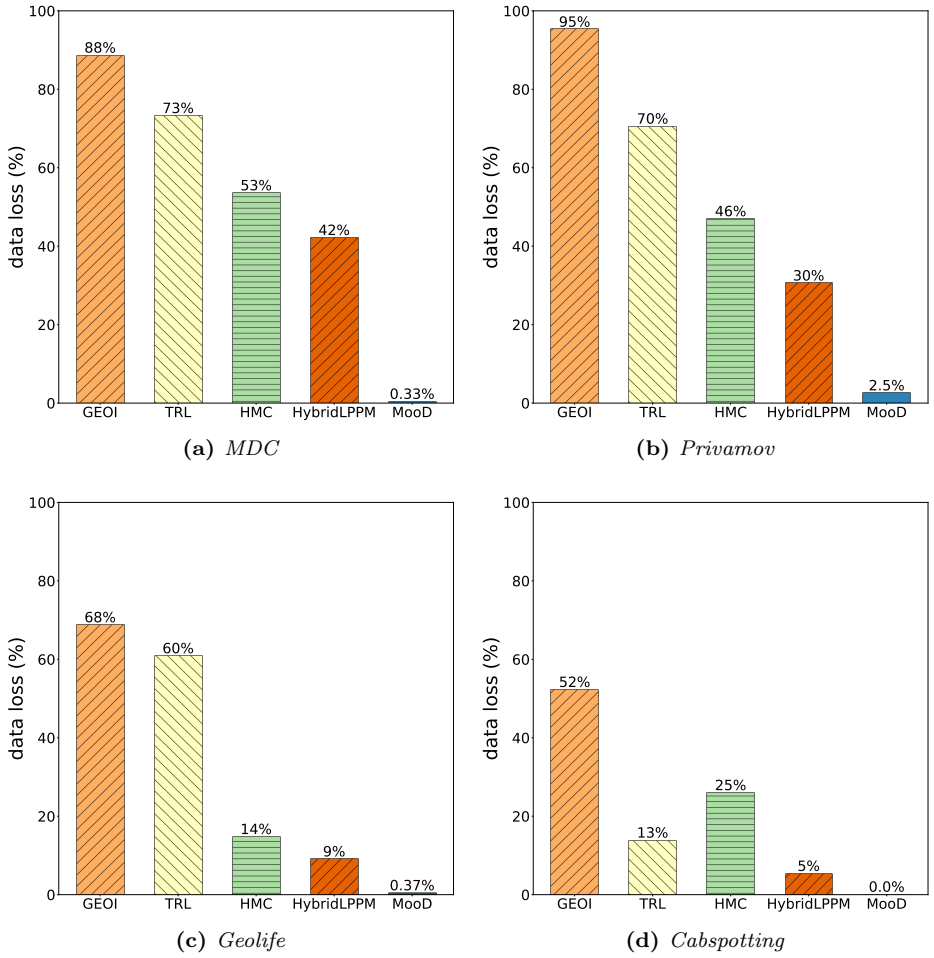


**(a)** *MDC* **(b)** *Privamov*

**(c)** *Geolife* **(d)** *Cabspotting*

**Figure 3.9:** *Data loss of* MOOD *vs. competitors.*

## 3.6 Summary

In this chapter, we presented MOOD, a centralized approach for location privacy protection. It is a user-centric multi-LPPM fine-grained protection system. Its main objective is to protect all users mobility data and in particular the minority of *orphan* users who are not protected by any single LPPM. MOOD can be used either for data publishing use case or for a crowd sensing campaign. It consists of three main components: The multi-LPPM composition search component which combines multiple LPPMs. The output of the first LPPM becomes the input of the following one, and so on until it finds a composition of LPPMs that succeeds the re-identification risk assessment. If the mobility trace is still vulnerable to re-

identification, the fine-grained protection component splits the mobility trace into sub-traces. We argue that short mobility traces may contain less discriminating information than longer ones. Therefore, re-identification attacks which are based on profiling will be less successful at re-identifying users because the discriminating mobility patterns extracted from the mobility trace of the user are separated and the sub-traces are assigned different IDs which makes them unlinkable to the original user. Finally, the best LPPM selection component chooses the candidate (single LPPM or multi-LPPM) that protects mobility data against re-identification while maximizing its utility. In our evaluation, we used the spatio-temporal distortion metric to measure the utility of the generated data.

In this chapter, we conducted experiments on four real-world mobility datasets to evaluate the effectiveness of MOOD. The results show that the proposed system is resilient to multiple re-identification attacks and can achieve a high level of privacy protection while maintaining an acceptable utility level and a low data loss. However, the main inconvenient of MOOD is that it assumes a *trusted proxy server* as the existing re-identification attacks require a centralized knowledge of users mobility data to construct mobility profiles. This is considered as a single point of failure because if the server is compromised, all users' mobility data is leaked and thus users' privacy can be violated. In the next chapter, we eliminate the assumption of the *trusted proxy server* which has access to raw mobility data and propose a novel location privacy risk assessment model to evaluate the privacy risks of sharing mobility data following a federated learning approach without accessing the raw mobility data.

# Chapter 4

# SAFER: a Federated Approach for Location Privacy Risk Assessment

## Contents

## 4.1 Motivation

In this chapter, we are interested in carrying out a privacy risk assessment for mobility data. It has been widely encouraged by the legislation (*e.g.*, Article 35 of the EU GDPR regarding risk assessment), and recommended by standardization institutes (*e.g.*, the NIST recommends to assess privacy risks for individuals arising from processing their data [114]). In this work, we present the concept of *uniqueness* as a means of location privacy risk assessment. It has been first introduced by De Montjoye *et al.* [64] on a dataset involving more than a million users. They demonstrate that only four spatio-temporal points, randomly drawn from a mobility trace, are enough to uniquely identify the originating user of the data in 95% of the cases. This study was performed on a CDR dataset where a subset of 2,500 mobility traces was used to compute the uniqueness. The same metric was used in [221] on the mobility data of more than half a million users to evaluate the effectiveness of anonymization techniques. In the same direction, another study has been conducted on smaller datasets over multiple types of sensors available in a smartphone (*e.g.*, GPS, WiFi, cellular data) [43]. While the above studies help raising the general public awareness about the sensitivity of mobility data, they can not be used as privacy risk assessment metrics as they require centralizing raw mobility data on a trusted server to analyze them. This raises serious privacy violation and data leakage risks, if the entity holding the whole raw data is compromised [6, 8, 3]. In this context, it is becoming increasingly important to devise mechanisms that avoid centralizing raw data in remote infrastructures.

In this chapter, we propose SAFER, the first solution for assessing the uniqueness of mobility data while keeping the data on the user's premises. To reach this objective, we model the uniqueness assessment problem as a machine learning classification problem and implement it using the FL paradigm where data remains on the user side. Specifically, SAFER consists of two main components: (i) an identity classification model based on mobility data and trained using the FL principles, and (ii) a local uniqueness evaluation component that is based on the computation of user anonymity sets.

SAFER's identity classifier is locally trained by mobile users before being aggregated by the FL server to build a global classification model that aims at identifying

to which user, a mobility data is likely to belong. In order to assess the uniqueness of mobility data for a given user, SAFER's FL classifier is used to produce a confidence vector with probabilities for the considered mobility data to belong to each class (*i.e.*, user). Using this vector, SAFER builds an anonymity set of users, namely a set of users whose probability of owning the mobility data is sufficiently close to the probability of owning the mobility data by the originating user. Thus, SAFER is able to determine that a given mobility data is unique if its anonymity set consists of a single element. Consequently, SAFER warns the user of a high privacy risk if that data is shared. Furthermore, besides the binary characterization of mobility data uniqueness as commonly performed in state-of-the-art solutions [64], SAFER is able to further analyze privacy risks. For that purpose, both the size and the entropy of the anonymity set of mobility data are considered. The higher the size and the entropy of the anonymity set, the lower the privacy risk of sharing that data.

We evaluate SAFER on four real-world mobility datasets [135, 169, 257, 1], with up to 10,000 mobile users and consider various representations of mobility data, namely GPS locations, POIs and trajectories. In order to assess the accuracy of SAFER, we compare its uniqueness estimation with a well-established centralized baseline [64]. Our experiments show that (i) SAFER estimates the uniqueness similarly to the state-of-the-art centralized solution while keeping users data private on their devices; (ii) SAFER's FL identity classifier is able to efficiently scale to thousands of users and (iii) SAFER is able to determine not only *if* but *how much* a mobility data is indistinguishable from other users' data. This last point is complemented by an evaluation of the privacy exposure of mobility data where a state-of-the-art re-identification attack (*i.e.*, AP-Attack [152]) is carried out to assess whether data that is considered as unique by SAFER is more subject to re-identification. Results show a clear correlation between the uniqueness estimated by SAFER and the re-identification attack success rate.

**Roadmap** The remainder of the chapter is organized as follows. We first describe the design principles of SAFER in Section 4.2. Then, we dive to a detailed description of SAFER in Section 4.3. Afterwards, we present our extensive experimental results in Section 4.4. Finally, we discuss the privacy limitations of SAFER and possible countermeasures in Section 4.5 and we draw our conclusions in Section 4.6.

## 4.2 SAFER Design Principles

In the following, we start by recalling some preliminary definitions in Section 4.2.1. Then, we describe the system and threat model in Section 4.2.2 and provide an overview of SAFER in Section 4.2.3.

### 4.2.1 Preliminary Definitions

***Mobility Data.*** Mobility data can be a single timestamped mobility point or a sequence of multiple mobility points forming a trajectory. The data point may correspond to the actual GPS location of a user, to the location of a cell network tower extracted from a CDR dataset, or to a POI extracted thanks to clustering algorithms [193] and represented by its centroid. A detailed description of mobility data is provided in Chapter 2, Section 2.1.1.

***Federated Learning.*** To recall, FL is a new machine learning paradigm that enables training machine learning models on data from different sources without the need to store the data at a central server. FL is performed in several learning rounds where set of users (known as workers) and a server (known as the federator) are involved. At the beginning, the FL server initiates the same model on all workers. For each FL round, the workers train locally the model with their private local data to improve the machine learning model and send the updated model to the federator. The latter then aggregates the received local models by averaging them, and produces a new version of the global model [159, 142, 112]. After the clients receive the aggregated model, a new FL round starts and the FL process is repeated until the aggregated global model converges. A background of the FL paradigm is provided in Chapter 2, Section 2.4.

### 4.2.2 System and Threat Model

Let $U = \{U_1, U_2, \ldots, U_N\}$ be a set of users of the system. Each user $U_i$ has its own mobility data stored locally on her device. This data can be collected at various times and using various sensors available on the user's device, depending on the actual

used mobile application. Specifically, in SAFER we consider three representations of mobility data: (i) mobility points corresponding to actual locations visited by a user, (ii) points-of-interest (POIs), and (iii) trajectories of a user. Using their mobility data, users participate in a federated training process (as defined in §4.2.1). The FL server is assumed to be honest-but-curious. It aggregates users model updates honestly without altering them but remains curious to inspect users updates in order to infer a training sample from a client's private dataset and thus can discover any private sensitive information about the user. To avoid this threat, we consider that users updates sent to the FL server are masked thanks to the secure aggregation protocol [39, 254]. Therefore, the FL server can only learn about aggregate client updates, which preserves the privacy of individual user contributions. In addition, the communication channels between the users and the server are encrypted. We assume that the training process is orchestrated periodically by the FL server (*e.g.*, each night) when users devices are available to perform computations (*i.e.*, they are likely to be idle, charging and connected to WiFi). These criteria are usually met at night and they are considered in many existing works [112, 40, 33]. With the advances of smartphone technologies, recent works discard these criteria and leverage the actual network connection of each device to provide a more dynamic federated learning protocol that evolves more rapidly and during the day [60]. These solutions are complementary to SAFER and could be considered in future work. We assume that users' devices are trusted. The case of malicious users is considered out of the scope of this chapter and is discussed in §4.5 where existing state-of-the-art countermeasures are listed. SAFER can be used by a user before sharing its mobility data, in an interactive way, to assess data uniqueness as further described in the following section.

### 4.2.3 SAFER Overview

SAFER (uniquenes*S* *A*ssessment with a *FE*derated lea*R*ning approach) is a user-side privacy risk assessment metric for mobility data. Its main objective is to locally measure, on users' devices, how unique is a given mobility data without accessing the other users' data. To reach this objective, SAFER operates in two distinct phases: a training phase and a uniqueness-evaluation phase as depicted in Figure 4.1. The training phase (depicted in the left part of Figure 4.1) is responsible

for training an identity classifier *ID Classifier* using the FL approach. During this phase, users exploit their local mobility data to train the *ID Classifier*, mask the model updates ($\Delta_i$) and send the masked gradients ($X_i$) to a FL server. The latter aggregates the received gradients without revealing users individual updates by following secure aggregation [39, 254] and sends back the resulting aggregate model to all the participants. This process is periodically repeated and is generally executed when user's devices are idle, charging and connected to WiFi (*e.g.*, at night time). The objective of the *ID Classifier* is to compute the probability that a given mobility data belongs to a given class (*i.e.*, a user). The *ID Classifier* is itself composed of two machine learning models that differ in the type of input data that is used (*i.e.*, mobility points or trajectories). Details about these two models are given in Section 4.3.1 and 4.3.2, respectively.

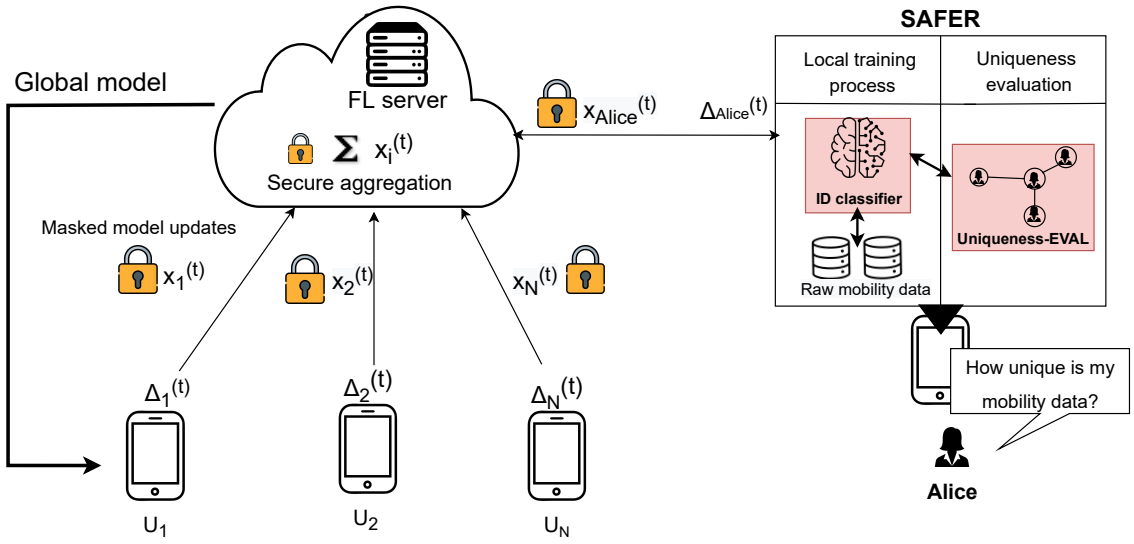


**Figure 4.1:** SAFER *overview*

Contrary to the training phase, which operates periodically on the background, the uniqueness-evaluation phase is an interactive phase that operates each time a user wants to assess the uniqueness of a given mobility data, say $I_P$ (right part of Figure 4.1). During this phase, the *Uniqueness-EVAL* component of SAFER uses the latest version of the *ID Classifier* to compute a confidence vector that contains the probability of $I_P$ belonging to each user of the system. Using this confidence vector, the *Uniqueness-EVAL* component creates an anonymity set, the size of which reflects how unique is $I_P$. Details about how the anonymity set is computed and how the uniqueness is inferred by *Uniqueness-EVAL* are further described in Section 4.3.3.

## 4.3 SAFER Detailed Description

In the following, we dive into the detailed description of the main components of SAFER as depicted in Figure 4.2. Particularly, we present DataPoint-CLASSIF for mobility data points (Section 4.3.1). Then, we propose Trajectory-CLASSIF as an alternative model to DataPoint-CLASSIF for trajectory data (Section 4.3.2), and finally we present Uniqueness-EVAL component (Section 4.3.3).

### 4.3.1 DataPoint-CLASSIF: Federated Identity Classifier for Data Points

The first type of classifier proposed by SAFER is called DataPoint-CLASSIF. It is a federated identity learning classifier for mobility data points, as depicted in the top left part of Figure 4.2. It takes as input a mobility point (*i.e.*, actual GPS location, a POI, or a CDR location) or a set of mobility points, and transforms them into a feature vector. This feature vector captures the spatial and temporal behavior of the user, in the form of a heat map to represent the mobility points. Precisely, a map is divided into cells of equal size. Each cell represents a spatial feature that has a value of either 1 or 0, depending on the presence or absence of the user at that location, respectively. In addition, DataPoint-CLASSIF considers a temporal feature to differentiate similar spatial mobility patterns occurring at different times, *e.g.*, a user visiting a museum where another user works. Here, the average hour of the day is used as temporal information. In the case of POIs, DataPoint-CLASSIF adds additional features which are the total number of actual points in a POI to represent its density and the POI duration to differentiate between short and long-length POIs.

Once the feature extraction step is prepared, DataPoint-CLASSIF uses a multi class logistic regression classifier (LR). We chose LR because it is a simple, yet effective model that can be trained locally on constrained devices. Periodically, when participants finish training their local model, they send their local model updates (*i.e.*, gradients) to the FL server. The latter aggregates users' model updates with the classical federated averaging method and produces an updated global model. This process is repeated over time, which allows DataPoint-CLASSIF to continuously learn new discriminating mobility patterns that make users more distinguishable

from each other. Finally, the classification model outputs a confidence vector that contains the probability for input data to belong to each user (*i.e.*, class) of the system. These probabilities reveal those users who hold similar data that allows SAFER to construct anonymity sets (§4.3.3).

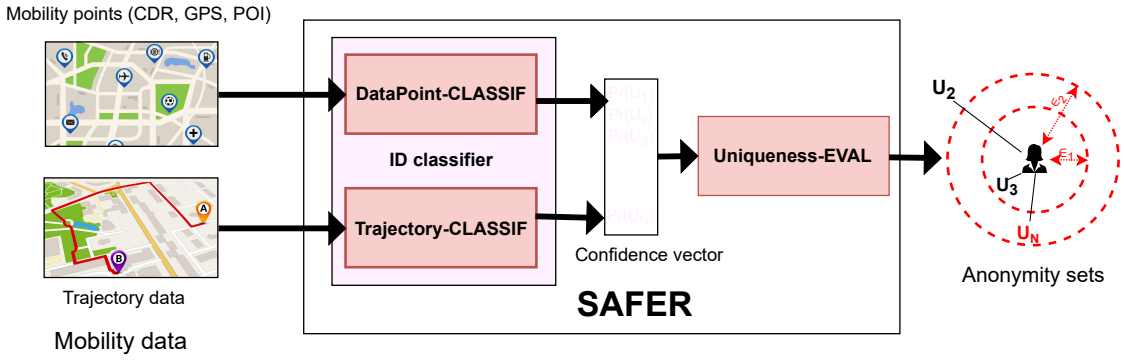


**Figure 4.2:** *Detailed description of* SAFER

### 4.3.2 Trajectory-CLASSIF: Federated Identity Classifier for Trajectory Data

In addition to individual mobility points or sets of individual mobility points, SAFER allows to evaluate the uniqueness of a trajectory as a whole. This can be considered in the context of data publishing use case where the user instead of assessing the uniqueness of each mobility point separately or by randomly choosing ones, the user evaluates the uniqueness of the whole trajectory at once before releasing it. To handle this complex data type, SAFER proposes Trajectory-CLASSIF, an alternative model to the DataPoint-CLASSIF previously presented in §4.3.1. This model relies on a recurrent neural network (RNN), and precisely on a bidirectional LSTM (Bi-LSTM) [116], which has been successfully used in the past to compute the similarity between mobility trajectories [90]. There are several reasons to use this architecture. First, it avoids extracting features that might lack relevant information. Second, it helps in processing long-term variable-length (*i.e.*, long mobility sequences). And last but not least, it captures recurrent movements and identifies the moving patterns thanks to the memory underlying LSTM (long short-term memory), *e.g.*, if a user visits locations corresponding to her home, work, then home again. Bi-LSTM is more complex than LR, however, its performance concerning trajectory data is

much better.

As depicted in Figure 4.2, Trajectory-CLASSIF takes as input a trajectory data. Each trajectory is transformed into a sequence of cells. A cell is a location with a unique cell identifier (*cellID*). Roughly speaking, a trajectory model aims at predicting a sequence of successive cells, namely a cellID that follows another cellID, *etc.*,. Here, we apply a simple yet effective technique inspired by natural language processing (NLP). Indeed, similarly to words in NLP, the frequency of locations follows a power law distribution [90]. For this reason, we embed each cellID using the word2vec technique to learn location associations [163]. We apply the common bag of words (CBOW) model where transformed trajectories are the context, and the cellID is the word to predict. Once the embedding is done, the latter produces for each cell, its feature vector that is used to feed the federated Bi-LSTM model. To link trajectories to their users, Trajectory-CLASSIF uses a dense output layer. The size of this layer corresponds to the number of classes in the system (*i.e.*, one output value for each class). Then the softmax function is used as an activation function [41] for multiclass classification. It provides probabilities of membership of an input trajectory to each class (*i.e.*, user identity label).

### 4.3.3 Uniqueness-EVAL: Anonymity Set Constructor

The last component of SAFER is Uniqueness-EVAL, as depicted in the right part of Figure 4.2. Uniqueness-EVAL first computes anonymity sets, then infers uniqueness based on the anonymity set size. To illustrate how Uniqueness-EVAL computes the uniqueness of mobility data, let $U = \{U_1, \ldots, U_N\}$ be the set of users participating in SAFER, $C_k$ the current version of the *ID Classifier*, $I_P$ a mobility data (mobility points or trajectory) belonging to a user $U_i \in U$. let $Pr(U_i)$ be the probability that $I_P$ belongs to the user $U_i$. This probability is computed by the classifier $C_k$. Note that the used *ID Classifier* depends on the type of $I_P$. Uniqueness-EVAL formally computes the uniqueness of $I_P$ as follows:

$$Uniqueness(I_P, C_k) = \begin{cases} 1, & \text{if } |AnonymitySet(I_P, C_k)| = 1 \\ 0, & \text{otherwise} \end{cases} \tag{4.1}$$

where

$$AnonymitySet(I_P, C_k) = \{U_j; |Pr(U_j) - Pr(U_i)| \leq \epsilon\} \quad (4.2)$$

In the above formula, the anonymity set $AnonymitySet(I_P, C_k)$, contains users $U_j$ with probabilities $Pr(U_j)$ equal to the probability of $I_P$ belonging to $U_i$ plus-minus a parametric $\epsilon$. These users hold mobility data that is spatially and temporally similar to $I_P$, such a similarity is detected thanks to the classifier $C_k$. Obviously, a user is unique if its anonymity set consists of a single element. Such users are more sensitive to privacy risks. In contrast, similar users are more likely to be in the same anonymity set and, thus, less distinguishable (*i.e.*, non-unique).

In this chapter, the uniqueness computation is based on anonymity set construction. Contrary to existing solutions that require sharing raw mobility data either with a trusted server (*i.e.*, centralized architectures) or with peers (*i.e.*, decentralized architectures), as discussed in Chapter 2, Section 2.3.2, SAFER does not require users to share their data. Instead, only vectors of probabilities produced by a federated learning model are used to form anonymity sets based on historical trained data, which is a novel way to compute uniqueness.

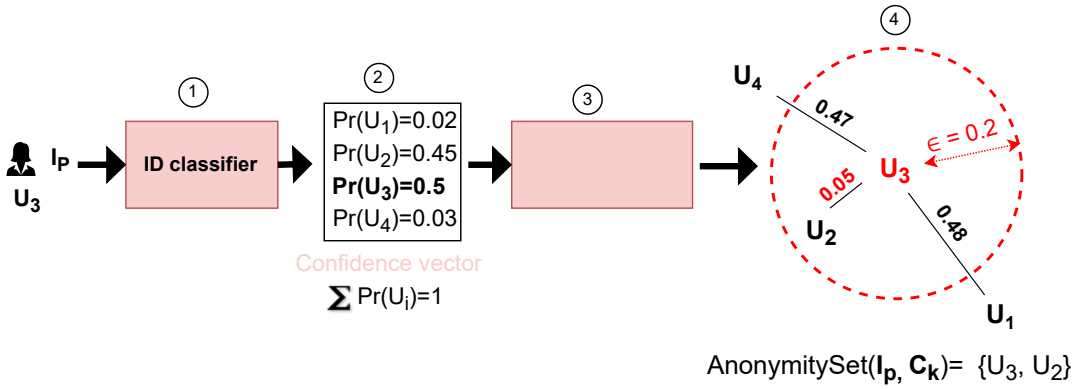


**Figure 4.3:** *Example of anonymity set construction in* SAFER

Figure 4.3 illustrates an example of how the uniqueness is computed in SAFER, where 4 users $U_1$, $U_2$, $U_3$ and $U_4$ are involved. Here, user $U_3$ wants to assess the uniqueness of a set of mobility points $I_P$ (depicted in the left part of the figure). To this end, $U_3$ uses the latest version of the ID classifier (step ① in the figure), which produces a confidence vector for $I_P$, $(0.02, 0.45, 0.5, 0.03)$ (step ② in the figure). In this vector, $Pr(U_i)$ represents the probability that $I_P$ belongs to the user $U_i$. From this figure, we observe that $U_3$ is the most likely owner of $I_P$ according to the classifier

$C_k$. Then, given $\epsilon$, SAFER evaluates the uniqueness in step ③ and constructs the corresponding anonymity set, *e.g.*, when $\epsilon = 0.2$, $AnonymitySet(I_P, C_k)$=$\{U_3, U_2\}$ (step ④ in the figure). Consequently, $I_P$ is not unique and it has an anonymity set of size two containing both $U_3$ and $U_2$. Intuitively, this means that $I_P$ is likely to be re-associated to either $U_3$ or $U_2$ with the same probability plus-minus 0.2.

## 4.4 Experimental Evaluation

In the following, we first present the implementation details of SAFER and our experimental environment in §4.4.1. Then, we present the evaluation of SAFER, which aims at answering the following questions:

- What is the impact of SAFER configuration parameters on the uniqueness of mobility data? (§4.4.2)
- How does the uniqueness measured by SAFER compare to a centralized baseline? (§4.4.3)
- Can we use SAFER dynamically? (§4.4.4)
- What can SAFER say about privacy risks beyond the binary measure of uniqueness? (§4.4.5)
- What is the potential re-identification risk of unique mobility data? (§4.4.6)
- Does SAFER scale in the number of mobile users and what is its computational cost? (§4.4.7)

### 4.4.1 Implementation and Experimental Environment

#### SAFER Implementation

Our system is developed in Python using different libraries, mainly including Pytorch [10] and Keras [7] for implementing the LR and the Bi-LSTM models, respectively. We use the S2Geometry library [11] for the conversion of mobility points into

cells of approximately equal size ranging between 100 meters and 200 meters. Additional libraries are also used for data preprocessing (*e.g.*, Pandas, Numpy, Word2Vec). In SAFER, we include participating users by considering the data coming from real-world mobility datasets.

In our implementation, we experimented different machine learning models for mobility data points including logistic regression (LR) [130], decision trees [82], and random forests [45]. The LR model was empirically chosen. It was a simple, yet effective choice in our work. In the FL protocol, the training process is performed by rounds. A round is a learning cycle where participating users train their model locally with their local data. In datasets with 30 days like MDC, Privamov, and Geolife, the training process is done chronologically every 24 hours (*i.e.*, 30 rounds), and in the MCOC-Shanghai dataset with 11 days of mobility, the training process is done every 6 hours (*i.e.*, 44 rounds). The uniqueness evaluation can be performed at any point in time using the latest version of the global model. For instance, if a user wants to test a given mobility data at round $i$, the user retrieves the latest generated model $C_{i-1}$ to assess the uniqueness. For both the LR and the Bi-LSTM model, we used a stochastic gradient descent optimizer with a learning rate ranging between 0.001 and 0.01 and a batch size equal to one. We considered the categorical cross-entropy loss function. In the Bi-LSTM model, the trajectories are embedded into feature vectors of size 300 with a window of size 50. We set the number of hidden layers to 250 layers. These values are empirically found after several experiments. To ease the reproducibility of our results, our system is available as an open source[1].

### Experimental Setup

Our experiments are carried out on a well-known academic cluster of machines [24] using machines with 2 CPUs intel Xeon E5-2650 v4, 12 core/CPU, and 128GB RAM. To accelerate the training process of our federated *ID Classifier*, the machines were equipped with a GPU of type GeForce RTX 2080 Ti. In all the experiments, we set the parameter $\epsilon$ to 0.001 unless specified differently. In addition, in our evaluation scenario, we vary $P$ *i.e.*, the size of the tested set containing mobility points between 1 and 5 for both SAFER and the baseline system. The latter is set to 3 when the focus is not on $P$ unless specified differently. The number of draws in both SAFER and

[1] https://gitlab.liris.cnrs.fr/bkhalfoun/safer (Private Access)

the baseline system is set to 100 draws (*i.e.*, for each user trace, we select randomly $P$ points 100 times). The length of tested trajectories for trajectory-based uniqueness is set to 6, 12, and 24 hours.

### Datasets

In our experiments, we use four real-world mobility datasets, namely Geolife [257], MDC [135], Privamov [169], and MCOC-Shanghai dataset collected from a major cellular operator in China in the city of Shanghai [1]. The latter contains call detail records of users initiating or receiving calls, messages, or any data. In all the datasets, each user has one mobility trace. In our experiments, we select only the most active month (*i.e.*, 30 days) of Geolife, MDC, and Privamov datasets and sample them with a resolution of 1 hour (*i.e.*, select a mobility point each 1 hour). Due to the large volume of the MCOC-Shanghai dataset, we processed it as follows: we removed redundant GPS coordinates and kept only users with more than 20 records. After that, we selected the most active users in the dataset. Details about each dataset's characteristics are described in Table 4.1. For the POI-based data representation, we used an existing spatio-temporal clustering algorithm to extract POIs [193]. It has two parameters: the diameter of the clustering area, and the minimum time spent inside a POI. They are respectively set to 200 meters and 15 minutes.

| Dataset | Type of data | Location | Area (km$^2$) | #Users/Traces | #POIs | #Records |
|---|---|---|---|---|---|---|
| MDC [135] | GPS | Lausanne | 41 | 144 | 4950 | 13788 |
| Privamov [169] | GPS | Lyon | 48 | 48 | 3123 | 4588 |
| Geolife [257] | GPS | Beijing | 16,411 | 42 | 1458 | 3084 |
| MCOC-Shanghai [1] | CDR | Shanghai | 6,340 | 10,000 | N/A | 846,239 |

**Table 4.1:** *GPS and CDR datasets*

### Baseline System

In order to assess the accuracy of the uniqueness measurement performed by SAFER, we use the well-established centralized solution proposed in [64] as a ground truth baseline. In this baseline system, the uniqueness of mobility data is measured in the following way. For a given dataset $D$ containing one mobility trace $T$ for each user, a random set of points $I_P$ picked from a trace $T$ is considered as unique in

$D$ (*i.e.*, $unique_{baseline}(I_P)$=1), if there is exactly one trace in $D$ that contains $I_P$. If another mobility trace $T'$ that has a set of points similar to $I_P$ is found, then, $I_P$ is considered non-unique. In this definition, the notion of similarity between mobility points is defined with respect to a spatial resolution $R_{spatial}$ and a temporal resolution $R_{temporal}$. $R_{spatial}$ is a distance under which two data points are considered as geographically similar while $R_{temporal}$ is a time difference under which two data points are considered as temporally similar. In our experiments, as we want to evaluate the uniqueness of a whole dataset with respect to a set size of $P$ points (where $P$ is a parameter), the above operation (*i.e.*, $unique_{baseline}(I_P)$) is done for all the traces of the dataset by randomly picking sets of $P$ mobility points in each mobility trace and repeating this process a given number of iterations. The number of iterations is set to 100 iterations (*i.e.*, draws). In addition, the spatial and the temporal resolution are set to 200 meters and 2 hours respectively.

### Evaluation Metrics

In this section, we formally define the metrics used to evaluate our system.

***Uniqueness rate***. The uniqueness rate of a dataset $D$ is computed, as described in Algorithm 2. There are two additional parameters that are used in the computation of the uniqueness rate of $D$ that are $P$, the number of mobility points on which the uniqueness is computed and $x$ the number of iterations. Specifically, for each mobility trace $T$ belonging to a user in $D$ (line 3), $x$ iterations are performed (line 4). In each iteration, $P$ points forming a set $I_P$ are randomly extracted from $T$ (line 5) and their uniqueness is computed (line 6). The uniqueness rate of $D$ is increased each time a random pick $I_P$ is identified as unique (line 7).

**SAFER *Accuracy***. The accuracy of SAFER is the ratio between the correct predictions of SAFER while considering the baseline system as ground truth and the total set of tested mobility data. The good predictions represent true positives (TP) and true negatives (TN) *i.e.*, what SAFER predicts the same as the baseline system. The accuracy of SAFER is defined as in Equation 4.3.

$$Accuracy_{\text{SAFER}} = \frac{TP + TN}{|D| * x} \tag{4.3}$$

***Classifier accuracy***. This metric measures the accuracy of the ID classifier.

**Algorithm 2** Uniqueness rate of mobility data

```
1: function UniquenessRate(D, P, x )
2:     cpt = 0                                   initialize a counter of unique data
3:     for each user's trace T in D do
4:         for each iteration it in range(x) do
5:             I_P= GenRandom(T, P) generate randomly P mobility points from T.
6:             if uniqueness(I_P, C_k)=1 then
7:                 cpt = cpt + 1
8:             end if
9:         end for
10:    end for
11:    rate = cpt / (|D|*x)
12:    Return rate
13: end function
```

Specifically, it measures the amount of data that the ID classifier can re-associate to its originating user (*i.e.*, when the originating user class got the highest probability of owning the data, *i.e.*, top 1). More formally, let $D$ be a mobility dataset, $I_P$ the randomly selected $P$ mobility points, $C_k$ the current version of the ID classifier and the id function, an oracle capable of revealing for each anonymous mobility data $I_P$, its owner identity. The classifier accuracy is defined as in Equation 4.4.

$$Accuracy_{\ IDClassifier} = \frac{\sum_{I_P \in D} \begin{cases} 1, & \text{if } C_k(I_P) = id(I_P) \\ 0, & \text{otherwise} \end{cases}}{|D| * x} \tag{4.4}$$

### 4.4.2 Preliminary Evaluation of SAFER

Since SAFER comes with a set of configuration parameters, one could ask what is the impact of these parameters on the uniqueness of mobility data. To answer this question, we first study the impact of $P$ on the uniqueness metric while $\epsilon$ is set to 0.1. Then, we study the impact of $\epsilon$ on the uniqueness metric while $P$ is set to 2. We also evaluate the accuracy of the ID classifier to assess its effectiveness. All the experiences are conducted on raw GPS data of MDC, Privamov and Geolife datasets.

### Impact of the number of points $P$

We depict in Figure 4.4 the uniqueness rate of SAFER while increasing the number of points P. Over all the datasets, we observe that the higher the value of $P$, the higher the uniqueness of mobility data. Specifically, in the MDC dataset (Figure 4.4a), the uniqueness considerably increases from 41% to 74% when $P$ varies between 1 and 5. This illustrates that users' mobility data becomes more distinguishable and thus more unique when $P$ is higher, which may lead to a higher privacy exposure *e.g.*, unique mobility data is more subject to re-identification. The same observation can be made in the Privamov dataset (Figure 4.4b). In particular, the uniqueness rate increases with +16% when $P$ ranges between 1 and 5. Finally, in the Geolife dataset (Figure 4.4c), the uniqueness is already high (82%) when $P$=1 and reaching up to 92% when $P$=3. This is due to the nature of the dataset, which contains 42 users, distributed over a large surface (16,411 $km^2$). As a result, each user's mobility behavior is distinct, and the visited locations are relatively sparse and thus more distinguishable.

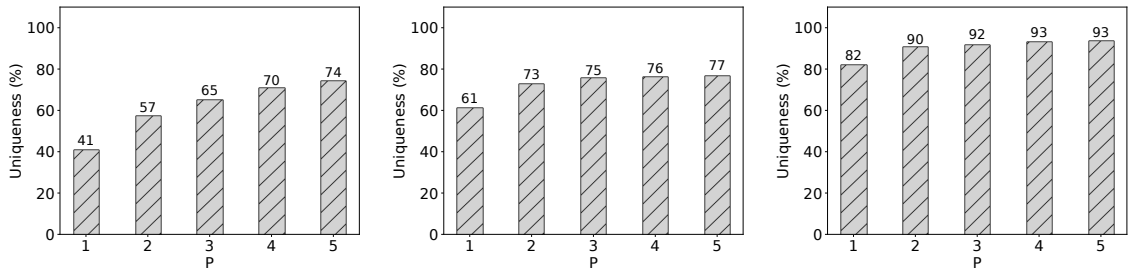


**(a)** *MDC* **(b)** *Privamov* **(c)** *Geolife*

**Figure 4.4:** *Impact of $P$ on uniqueness results*

### Impact of $\epsilon$ Parameter

In Figure 4.5, we illustrate the evolution of the uniqueness rate while varying the parameter $\epsilon$. The results show a similar tendency to the parameter $P$. In particular, in the MDC dataset (Figure 4.5a), the uniqueness rate increases considerably (+17%) while reducing the $\epsilon$ value from 0.1 to 0.0001. Decreasing $\epsilon$ leads to minimizing the distance between the probability of being the originating user of the mobility data and the probability of owning the data by others in the system. In addition,

the number of users influences the impact of $\epsilon$. Specifically, MDC has a higher number of users in comparison to Privamov and Geolife datasets (*i.e.*, 144 users *vs.* 48 and 42 users), distributed over a small area of 41 $km^2$. Thus, it is more likely that the remaining 43% of non-unique mobility data (around 6192 mobility records when $\epsilon$=0.1) become unique when $\epsilon$ is set lower. Whereas in the Privamov and Geolife datasets (Figure 4.5b and Figure 4.5c), the uniqueness rate slightly increases (+9% and +7% respectively). This is due to the lower number of users (48 and 42 respectively), distributed over a relatively large area (*i.e.*, 48 $km^2$ and 16,411 $km^2$). Thus, when reducing the value of $\epsilon$, it is less likely to create singleton anonymity sets, as users have sparse locations. Specifically, there are only 27% and 10% (in comparison to 43% in MDC) of non-unique mobility data for the Privamov and Geolife datasets (*i.e.*, around 1296 and 420 records, respectively).

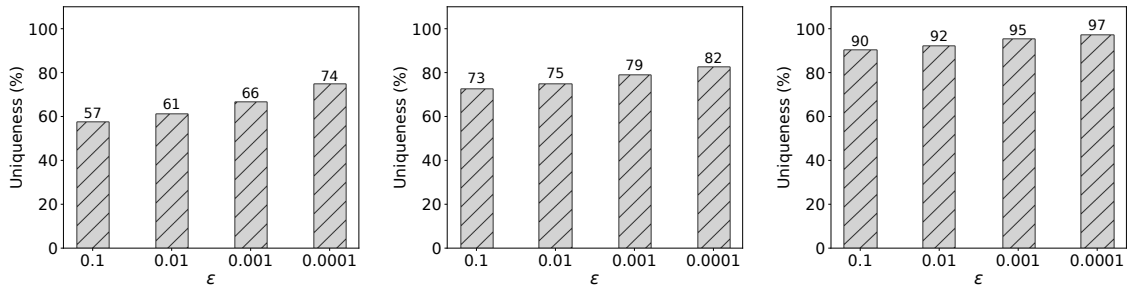


**(a)** *MDC* **(b)** *Privamov* **(c)** *Geolife*

**Figure 4.5:** *Impact of $\epsilon$ on uniqueness results*

### Accuracy of the ID Classifier

As SAFER is based on an ID classifier, we evaluate how the latter is accurate using the metric defined in Section 4.4.1. Towards that purpose, we consider the latest FL model of the ID classifier and the same tested data for $P$ ranging from 1 to 5 as in the previous experiments. Figure 4.6 illustrates the accuracy of the ID classifier on raw GPS data taken from MDC, Privamov, and Geolife datasets. Specifically, for the MDC dataset, the accuracy increases from 41% to 72% while increasing the size of $P$, which means that the ID classifier can learn discriminative and distinguishable mobility patterns from one user to another. Increasing $P$ leads to higher discrimination between users. In the Privamov dataset, the accuracy increases slightly from 57% to 64%, and finally, in the Geolife dataset, we record a higher

accuracy reaching up to 93%, which is due to the nature of the dataset. These results confirm that our approach relies on an effective ID classifier.

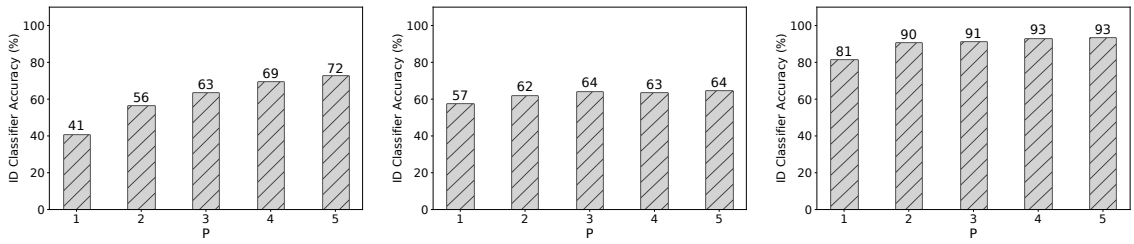


**(a)** *MDC* **(b)** *Privamov* **(c)** *Geolife*

**Figure 4.6:** *The accuracy of the ID classifier of* SAFER

### 4.4.3 SAFER Accuracy

In this section, we evaluate the uniqueness rate of SAFER in comparison to a centralized well-established baseline system in Figure 4.7. Then, we measure the accuracy of SAFER while considering the baseline as ground truth in Table 4.2. We conducted these experiments on raw GPS data of MDC, Privamov, and Geolife datasets. Further, we analyze the uniqueness of different mobility data representations taken from the same datasets using SAFER. To recall, we set the value of $\epsilon$ to 0.001.

#### SAFER *vs.* the Baseline System

In this experiment, we compare SAFER and the baseline system in terms of uniqueness rate using the same randomly selected mobility points for both systems. In addition, for a fair comparison between the latter two systems in terms of mobility background knowledge, we consider the latest ID classifier of SAFER, which has been trained on the whole users' mobility dataset after multiple rounds, as for the baseline system, which already knows the raw mobility dataset. Results are depicted in Figure 4.7. They show that SAFER behaves similarly to the baseline system when increasing $P$ except at $P$=1, SAFER provides a higher uniqueness rate. Specifically, in the MDC dataset (Figure 4.7a), the uniqueness rate increases for both systems with an average negative gap of -9% for SAFER from $P$=3 to $P$=5. In contrast, we notice that for $P$=1, SAFER exceeds the baseline with +18%. This behavior

is acceptable because the ID classifier learns the presence or absence of a user at different locations and times. Therefore, the classifier knows where the user mobility is concentrated over time. As a result, SAFER can distinguish unique mobility data earlier. In the Privamov dataset (Figure 4.7b), the evolution of the uniqueness between the baseline and our system is relatively different until $P$=3. More precisely, in the beginning, SAFER provides a higher uniqueness rate than the baseline (77% *vs.* 52%). This is for the same reasons mentioned for the MDC dataset. After $P$=3, both systems yield almost the same result. In the Geolife dataset (Figure 4.7c), both SAFER and the baseline provide a high uniqueness rate whatever the value of $P$. This behavior is due to the nature of the dataset, where the distribution of a small number of users' mobility points over a large area makes the users' mobility data more unique and thus easily distinguishable.

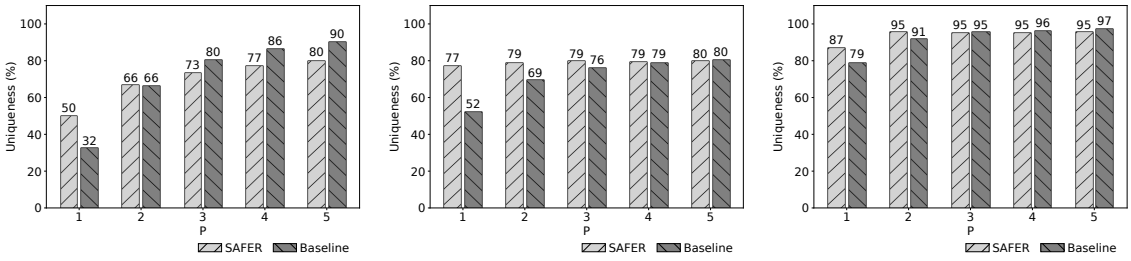


**(a)** *MDC* **(b)** *Privamov* **(c)** *Geolife*

**Figure 4.7:** *Uniqueness with* SAFER *vs. the baseline system*

***Observation 1 :*** SAFER *provides similar results of uniqueness compared to the well-established centralized system for computing uniqueness of mobility data [64]*

### Accuracy of SAFER

To better compare the uniqueness measurements of SAFER with the baseline system, we compute the accuracy, previously defined in Section 4.4.1. We consider the baseline system results as a source of *ground truth* for our comparison. The results are shown in Table 4.2. Overall datasets, the accuracy increases when $P$ is high. Specifically, in the MDC dataset, SAFER already achieves a high accuracy (75%) when $P$=3 or higher. It means that SAFER tends to produce similar predictions to the baseline system. This is because SAFER can approach the baseline system thanks to its learning paradigm. It builds a global knowledge about users' mobility that helps in

uniquely characterizing users without accessing their raw data. Similar behavior is observed in the Privamov dataset with 81% of accuracy when $P$=3. In the Geolife dataset, the results are high reaching up to 94% accuracy when $P$=3. As previously explained, this is because users' mobility data is naturally distinguishable.

***Observation 2 :*** *From $P = 3$,* SAFER *provides accurate results in comparison to the baseline as a source of ground truth.*

| Dataset | MDC | | | | | Privamov | | | | | Geolife | | | | |
|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|
| $P$ | 1 | 2 | 3 | 4 | 5 | 1 | 2 | 3 | 4 | 5 | 1 | 2 | 3 | 4 | 5 |
| *Accuracy* $_{\text{SAFER}}$ (%) | 59% | 66% | **75%** | 79% | 82% | 64% | 78% | **81%** | 83% | 84% | 79% | 91% | **94%** | 95% | 96% |

**Table 4.2:** *Accuracy of* SAFER

### Analyzing Different Mobility Data Representations with SAFER

In the following, our objective is to evaluate the uniqueness rate of different mobility data representations using SAFER. Specifically, we measure uniqueness for (i) raw GPS data and POI-based data, then (ii) trajectory-based data with different lengths of trajectories (6 hours, 12 hours, and 24 hours). Results are shown in Figure 4.8. Overall datasets, we observe that POI-based data is more unique than raw GPS data when $P$=3 and $\epsilon$=0.001. Specifically, 81%, 82%, and 95% of POIs are unique in MDC, Privamov, and Geolife, respectively. This is because POIs are semantically sensitive locations such as home or workplace that uniquely identify a user and characterize their mobility behavior.

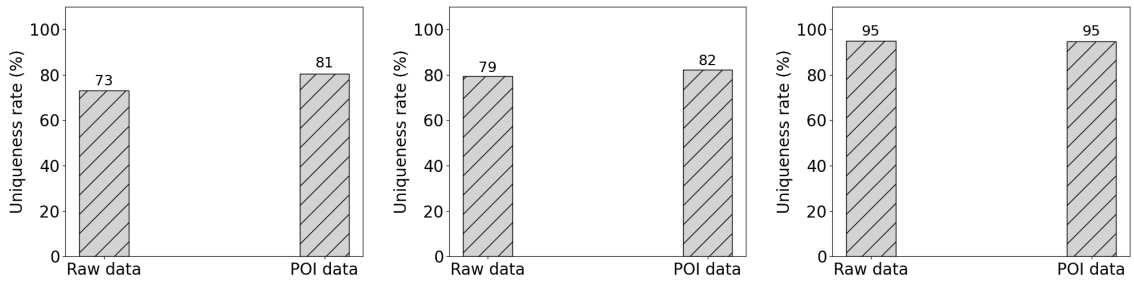


**(a)** *MDC* **(b)** *Privamov* **(c)** *Geolife*

**Figure 4.8:** SAFER *with mobility data points.*

***Observation 3 :*** *POIs are the most sensitive mobility data representation – even more sensitive than actual location data points –, they allow to better and uniquely distinguish users from each other.*

For trajectory-based data, results are depicted in Figure 4.9. In this experiment, we measure the uniqueness of all trajectories, where each of them is considered as a whole instead of using random data points. For that purpose, SAFER involves a specific ID model (*i.e.*, Trajectory-CLASSIF) to handle trajectories. The results show that the uniqueness rate increases when the trajectory length is longer (+9% in MDC and +20% in Geolife), which means that longer trajectories allow capturing recurrent mobility habits that uniquely identify the user. Unlike MDC and Geolife datasets, in the Privamov dataset, trajectories with a length of 12 hours provide a higher uniqueness rate than trajectories of 6 hours or 24 hours. This is because shorter traces (6 hours) contain less discriminative patterns in comparison to trajectories of 12 hours. However, trajectories with 24 hours length are less unique than the latter (-3%). This is because the Trajectory-CLASSIF model detects similar moving patterns that it could not detect on shorter traces. As a consequence, users' trajectories become non-unique.

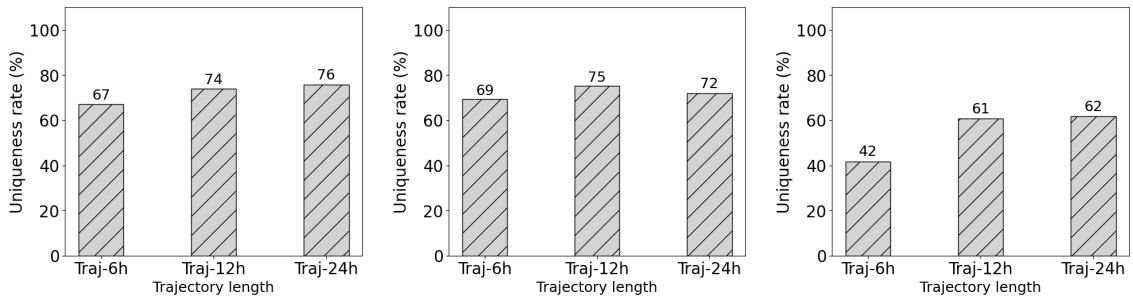


**(a)** *MDC* **(b)** *Privamov* **(c)** *Geolife*

**Figure 4.9:** SAFER *with trajectory data.*

### 4.4.4 SAFER in a Dynamic Use Case

In this section, our objective is to compute the uniqueness of mobility data points in a dynamic use case. A dynamic use case may correspond to scenarios where users want to share their POIs in real-time (*e.g.*, when a user visits a HIV center,

participates in a manifestation, or takes a drink at a gay bar). Before posting such locations, the user may wonder about their location privacy and thus perform a location privacy risk assessment. To this end, we experiment SAFER with $P$=1 and $\epsilon$=0.001 on POI-based data of the MDC, Privamov, and Geolife datasets. Results are depicted in Figure 4.10 and present the uniqueness rate per round. In each round, we compute the uniqueness of POIs extracted on that day using the ID model of the previous day (*i.e.*, the latest model in our configuration). Overall datasets, we observe a significant variation in the uniqueness rate between rounds with a tendency to increase. This is due to the spatio-temporal locality of the generated POIs. In other words, if the tested POIs are similar spatially/temporally to the past mobility of other users, this creates a negative slope, whereas when the tested data is different from the past trained data by the FL model, in this case, the uniqueness rate has a positive slope. In particular, the uniqueness rate varies between 14% - 69% for the MDC, 58% - 100 % for Privamov, and 6% - 80% for Geolife dataset.

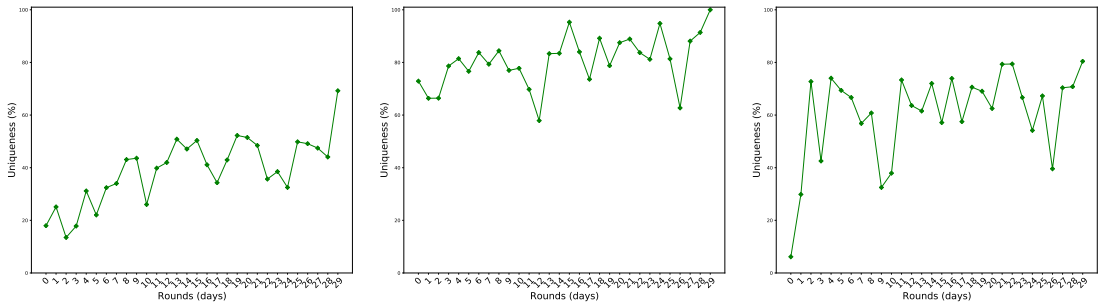


**(a)** *MDC* **(b)** *Privamov* **(c)** *Geolife*

**Figure 4.10:** *Uniqueness with* SAFER *over multiple rounds*

***Observation 4 :*** SAFER *can be used dynamically in real-time to evaluate the uniqueness of mobility data points.*

### 4.4.5 Beyond Uniqueness: Evaluation of Anonymity Sets

This section focuses on non-unique users' mobility data within each data representation. Precisely, we aim at measuring the size of anonymity sets for the remaining non-unique mobility data in raw GPS, POI-based, and trajectory-based data in MDC, Privamov, and Geolife datasets. We set $P$ to 3 and $\epsilon$ to 0.001. The results are illustrated in Figure 4.11 and present the size of anonymity sets in the form of

boxplots with a mean value represented with a red line. The results show that raw and POI-based non-unique data are blended with many other users when compared with trajectories. Specifically, the size of anonymity sets is almost equal to the total number of users in MDC, Privamov, and Geolife datasets with a mean value around 130, 47, and 37, respectively.

***Observation 5 :*** *Non-unique POIs and actual location data points are mostly hidden in large anonymity sets. These locations usually correspond to public locations that are frequently visited by most of the users.*

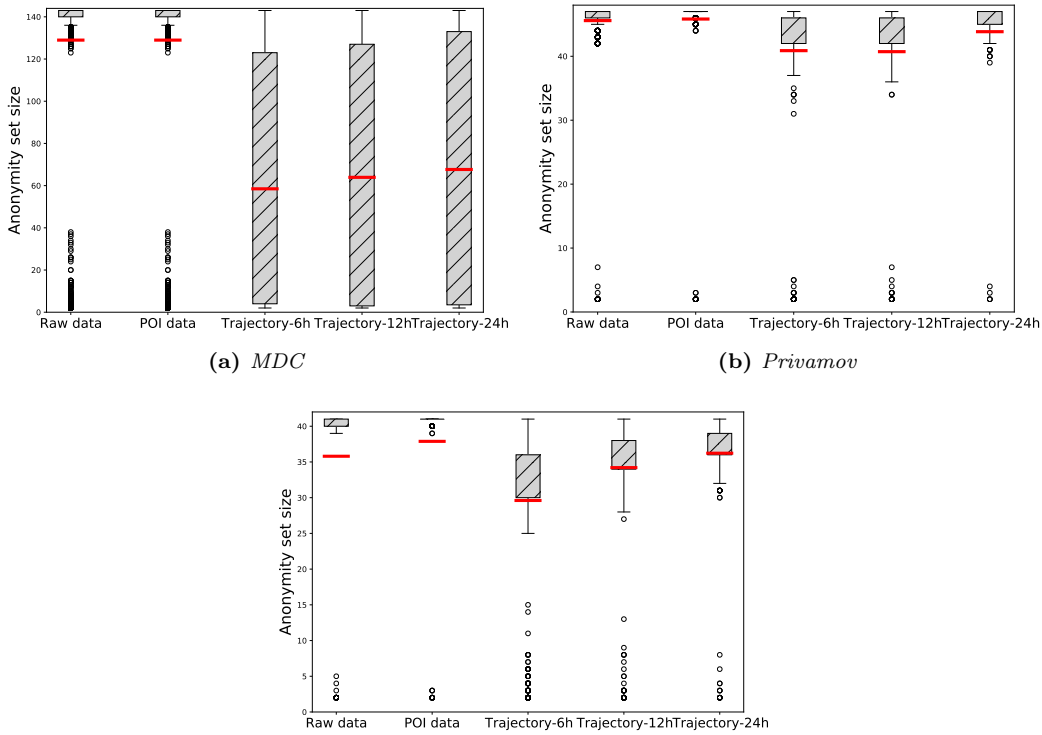


**(a)** *MDC* **(b)** *Privamov*

**(c)** *Geolife*

**Figure 4.11:** *Anonymity set size of non-unique data for different data representations*

For trajectories, the anonymity set size shows a high variability in the MDC dataset (Figure 4.11a). Specifically, it varies between 2 and 125 users on average, with a mean value ranging between 58 and 65 users in the anonymity set. Larger anonymity sets reflect a similar mobility behavior of humans (low privacy risk), for instance, when users borrow the same paths (*e.g.*, highway) to go to work. In contrast, the small anonymity sets reflect an uncommon mobility behavior (high

privacy risk), for example, going to the emergency center at night. In the Privamov and Geolife datasets (Figure 4.11c and Figure 4.11b), there is less variability in the anonymity set size. Particularly, the Privamov dataset represents the mobility of students, staff members, and their relatives from three universities in Lyon. The participants are distributed over a small area of 48 $km^2$. As a result, non-unique trajectories are more likely to be similar and thus grouped in the same anonymity set, *e.g.*, students living in the same residence and studying at the same university. For the Geolife dataset, non-unique trajectories are grouped in anonymity sets of size almost the total number of users in the dataset, which means that the selected value of the parameter $\epsilon$ is relatively optimal as it separates between unique and non-unique trajectories while there are only few outliers.

### 4.4.6 Evaluation of Privacy Exposure of Mobility Data

In this section, we study the potential re-identification risk of sharing mobility data if the latter falls between the hands of a malicious entity. To this end, we use an existing state-of-the-art re-identification attack, called AP-Attack [152]. This attack assumes a background knowledge for each user from which the attack constructs a mobility profile in the form of a heat map synthesizing the user's past mobility. Then, upon receiving anonymous mobility data, the attacker tries to guess to which user the data belongs. We use this attack to assess whether mobility data that is considered unique by SAFER is indeed re-identified by an adversary who runs this attack. The results are depicted in Figure 4.12. This figure illustrates both the re-identification rate obtained by the adversary and the uniqueness estimated by SAFER on mobility data taken from MDC, Privamov, and Geolife datasets when $\epsilon$=0.001. Specifically, the green bar represents unique and re-identified data, the red bar represents unique and not re-identified data, the blue bar represents non-unique and re-identified data, and finally, the orange bar represents non-unique and not re-identified data. From these results, we observe that the green color dominates in most datasets with values ranging between 33% - 63% for MDC, 57% - 76% for Privamov, and 77% - 90% for Geolife. In other words, a large portion of unique data by SAFER is re-identified by a potential adversary. For the unique and not re-identified mobility data (*i.e.*, red color), SAFER has a relatively conservative approach to assess the uniqueness of mobility data with a small portion of data. Specifically, we record an average of 21%,

15%, and 10% of data that falls in this case. Nevertheless, there is still a portion of data (blue color) that SAFER could not consider as unique but that an adversary was able to re-identify. This amount of data ranges between 4% and 14% on average overall the datasets.

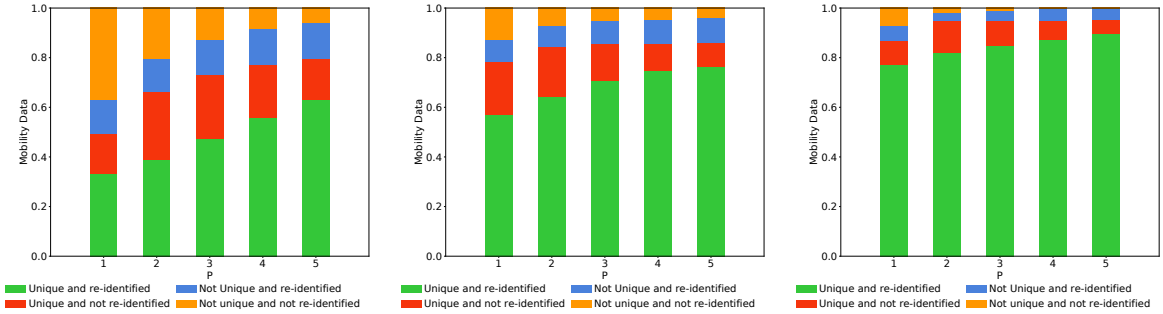


**(a)** *MDC* **(b)** *Privamov* **(c)** *Geolife*

**Figure 4.12:** *Re-identification and uniqueness correlation.*

### 4.4.7 SAFER Scalability and Computational Cost

In this section, we evaluate the scalability of SAFER when increasing the number of users and study its computational cost.

#### Scaling the Number of Users in SAFER

In this section, we evaluate the uniqueness rate while scaling the number of mobile users from 500 to 10,000 in the MCOC-Shanghai dataset. In addition, we compare our results to the baseline system when $P$= 3 and $\epsilon$= 0.0001. Results are depicted in Figure 4.13. They illustrate that increasing the number of users affects the uniqueness rate moderately. Specifically, for both SAFER and the baseline system, the uniqueness decreases by -17% and -15%, respectively. It means that a crowded place with the mobility of many users makes the user mobility more probable to be similar to the mobility of the others and thus less distinguishable. Despite this slight slope, the uniqueness is still high and close to the baseline with an average gap of 4%. This result demonstrates that our ID classifier is still effective when there is

a large number of classes (*i.e.*, users), and it is still able to discover discriminating mobility patterns, which uniquely identify users even in the crowd.

***Observation 6 :*** SAFER *and its federated ID classifier are able to handle a large number of users and their underlying classes.*

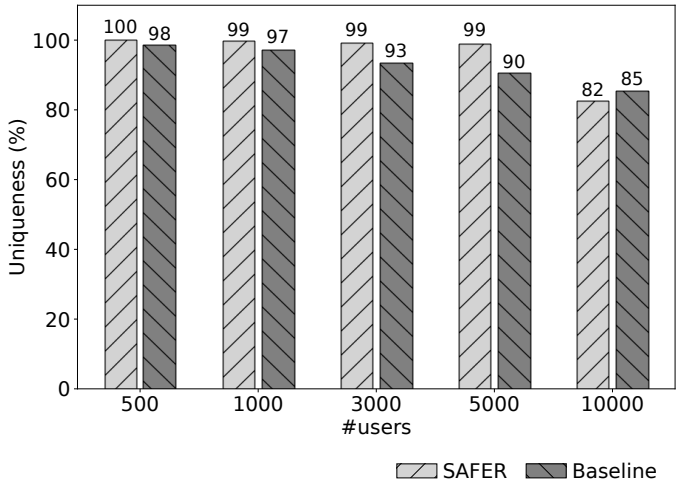


**Figure 4.13:** *Uniqueness with* SAFER *vs. baseline system at different scales*

### Computational Cost of SAFER

In this section, we want to answer the following question: "What is the computational cost of SAFER while scaling the number of mobile users?". To this end, we measure the run-time induced by both SAFER and the baseline system, considering 500, 1,000, 3,000, and 10,000 users of the MCOC-Shanghai dataset. In SAFER, the computation of uniqueness includes: (i) converting the data to feature vectors, (ii) loading the latest ID classifier and predicting the confidence vector of the mobility data, and finally, (iii) building the anonymity set. In the baseline system, we consider the scenario where the algorithm has to verify the mobility trace of each user in the dataset (*i.e.*, worst-case scenario). The results are illustrated in Figure 4.14a where we present the average execution time of uniqueness computation in both the baseline system and SAFER. The results show that SAFER has a considerably lower execution time (between 22 ms and 85 ms) in comparison to the baseline, which reaches up to 8,126 ms while increasing the number of users from 500 to 10,000. Thus, we notice that increasing the number of users by 20 times increases

the execution time by only 4 times, which is negligible in comparison to the baseline, which increases the execution time by up to 20 times. Therefore, SAFER is less costly because, instead of iterating over all users' data to evaluate uniqueness, our system uses the trained classifier to generate a confidence vector that constructs the anonymity set. The size of the confidence vector and the different operations mentioned above have a low influence on the execution time.

In addition, as SAFER follows the FL protocol, we measure the execution time of the ID classifier during a learning round per user. The results are depicted in Figure 4.14b. They show that the training time of a learning round per user increases considerably, from 467 ms to 3,861 ms, when the number of users scales in the system. This is due to the complexity of the model (higher number of classes and features). Our results are acceptable since we train the model periodically at night, when the user's device is idle, charging, and connected to WiFi. Furthermore, this experiment is conducted on a desktop machine (see Section 4.4.1), but it is still practical while using a smartphone. Indeed, there exist frameworks for on-device optimized model training which may help us in the future to deploy SAFER. Examples of existing FL systems include Apple's Core ML [2], the model used by Google's Gboard [13], or the one used for adaptive brightness on Google Pixel [5].

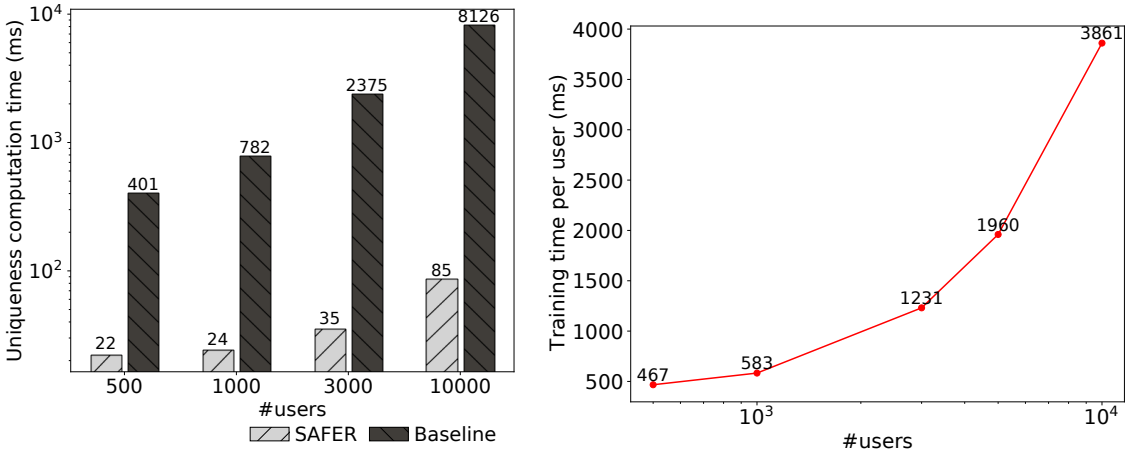


**(a)** *Uniqueness computation time of* SAFER *vs. baseline system*

**(b)** SAFER*'s model training time*

**Figure 4.14:** *Local computational cost of* SAFER *vs. baseline system*

## 4.5 Privacy Discussion

Although SAFER has promising results, we are aware that there is still room for improvement. We discuss in the following the privacy limitations of SAFER and possible countermeasures which could be combined with SAFER in the future.

### 4.5.1 Compromised Users in SAFER

In the FL protocol, participating users are considered as a critical component of the architecture. They can possibly act as passive or as active adversaries. Passive adversaries are honest-but-curious (*i.e.*, semi-honest), they try to infer private information about other users relying only on the aggregated model without deviating from the FL protocol. In this direction, inference attacks are classified into white-box and black-box attacks. Users have a full access to the FL model in the former and can only query it in the latter [26]. Unlike passive adversaries, active adversaries or malicious users try to learn private information of other honest users [252] and may deviate from the FL protocol. Specifically, malicious adversaries can either inject poisonous data to affect the aggregated model during the training phase [106, 215, 83, 226] or can send a malicious model instead of the trained model to the aggregator so that it outputs the target result desired by the adversary [37, 34, 23]. To mitigate such threats, many techniques are proposed in the literature. Some techniques aim at detecting poisoned inputs that deviate from benign inputs [83]. Others rely on clustering [218, 37] and anomaly detection [20] to capture malicious model updates. However such solutions require analyzing the FL client inputs in clear while SAFER is based on secure aggregation which aims to mask these inputs. Therefore, these solutions do not fit with SAFER. Alternative solutions are proposed to verify model updates while keeping them protected [177]. Moreover, homomorphic encryption can also be considered [124] in addition to masking-based secure aggregation where users are grouped into clusters and the aggregation is done per cluster [255, 231]. Finally, hardware solutions based on trusted execution environments (TEEs) are also solutions that are effective and complementary with SAFER. Indeed, TEEs can be used to protect the FL system in an end-to-end manner both on the client side using solutions such as DarkneTZ [165] or GradSec [161] and on the server side using solutions such as Sear [256]. In these solutions, all the

model parameters are encrypted and are manipulated in clear only inside hardware enclaves. These solutions also protect against poisoning attacks as the code can not be modified by adversaries even with root privileges (*e.g.*, the operating system).

### 4.5.2 Malicious Aggregator in SAFER

In this chapter, we consider the honest-but-curious threat model, which does not protect against a malicious aggregator that would tamper with the FL model to reveal sensitive information about a target user [184]. An attacker can also infer sensitive information by exploiting the joint model to run membership inference attacks or reconstruction attacks [176, 143]. To prevent such attacks, three categories of methods are proposed in the literature. Differential privacy (DP) is one of them [171, 16, 227, 214]. It adds noise to the local trained models before reaching the malicious aggregator. However this method affects the accuracy of the global model and thus can deteriorate the quality of the proposed service. Secondly, the secure aggregation protocol is also an established security primitive in FL. However, it is not sufficient as a malicious aggregator can select specific values for the aggregated model so that to disclose the input of a target user [184]. To avoid such a behavior, a verification of the aggregation result is required. Existing solutions are proposed in this direction for FL applications [109, 241, 253]. Specifically, the authors in [241, 253] use a homomorphic hash function (HHF) to verify the aggregation outcomes. However, it needs an increasing computation and communication overhead due to the increasing dimension of input data which is a clear limitation since the performance of the ML model highly depends on its size (*i.e.*, number of parameters). To handle this issue, the authors of VeriFL [109] focus on designing a verification scheme dedicated for secure aggregation of FL applications with high dimension inputs. They proposed a commitment/decommitment scheme to verify the aggregation combined with the HHF proposed in [29] to reduce the communication overhead. Finally, hardware solutions (TEEs) can also be used to protect against a malicious FL server [256, 166]. TEEs have the advantage of preserving the accuracy of the trained model and they are compatible with SAFER.

## 4.6 Summary

In this chapter, we presented SAFER, the first distributed user-side privacy risk assessment metric for mobility data. SAFER considers different types of mobility data (GPS or CDR), and different representations of data (actual locations, POIs, trajectories). Unlike state-of-the-art solutions, SAFER is able to determine not only *if* but *how much* a mobility data is indistinguishable from other users' data. To this end, it follows a federated learning approach, to locally and periodically train, on users' devices, a classification model that aims at identifying to which user a mobility data is likely to belong. The underlying model differs according to the input mobility data. For single data points, SAFER uses a simple, yet effective LR model, whereas for trajectories with sequential data points, SAFER uses a Bi-LSTM model to capture recurrent moving patterns. Both models allow building anonymity sets, which are used by SAFER to estimate the uniqueness of mobility data. Our experimental evaluation on four real-world mobility datasets show that SAFER is able to provide comparable results to a well-established centralized baseline and to efficiently scale to thousands of users. In addition, there is a clear correlation between the uniqueness estimated by SAFER and the re-identification success rate performed by a potential adversary.

# Chapter 5

# EDEN: Enforcing Location Privacy Protection through Location Privacy Risk Assessment: a Federated Learning Approach

## Contents

## 5.1 Motivation

In this chapter, we are interested in addressing the challenge of location privacy in the use case of crowd sensing applications. A crowd sensing application is an application where a set of (paid or volunteer) users carry a device equipped with a GPS and an environmental sensor (*e.g.*, an $NO_2$ sensor), which could be their own smartphone or a dedicated device. Along their journey, the application collects timestamped, geo-located traces with the corresponding environmental measurements (*e.g.*, pollution measurements). Then, it periodically sends this data to a central server called the Mobile Crowd sensing Server (MCS), which aggregates the collected data and provides updated information to its clients. The updated information can be related to traffic congestion monitoring [100], air quality monitoring [66] or radioactivity level monitoring near nuclear sites [46]. However, the downside of these applications is that the collected data may constitute a serious threat to the participants' privacy if this data falls between the hands of curious/malicious adversaries. Indeed, various studies have shown the privacy threats affecting mobility data. This has been extensively presented in Chapter 2 with a particular attention on re-identification attacks which jeopardize user's identity disclosure.

To overcome the above threats, the research community has been actively proposing diverse LPPMs offering different guarantees in terms of privacy, utility and performance [199, 21, 196, 119]. In this context, a problem that mobile app developers aiming at enforcing privacy-by-design have to solve is : "*how to objectively compare the privacy vs. utility tradeoff offered by different LPPMs and choose the right one ?*" For instance, how to decide whether an LPPM enforcing k-anonymity [212] (with a given value of $k$) is better than another one enforcing $\epsilon$-differential privacy [69] (with a given value of $\epsilon$)? To answer this question, the regulator (*e.g.*, the EU GDPR in article 35) requires to carry out privacy risk assessment, which in our context translates into assessing which solution yields the smallest re-identification risk.

In practice, solutions that have been explored in the literature to select among a set of LPPMs generally rely on re-identification attacks [127, 192, 206]. Specifically, these solutions apply various LPPMs on a given trace and choose the LPPM (and its corresponding configuration) that better resists a given set of re-identification attacks [87, 193, 152]. The role of these attacks is to link anonymous traces to past

user data. However, to reach this objective, the proposed solutions assume a trusted proxy server as existing re-identification attacks are centralized: they build user profiles using past unprotected mobility data and use them to estimate to whom a given protected trace belongs to.

In this chapter, we overcome this assumption and propose EDEN, the first effective privacy-preserving solution for mobility data that performs re-identification risk assessment without requiring to send raw data to a remote server and that provides high data utility. Specifically, EDEN operates in two phases: (1) a phase on which a re-identification risk assessment model called FURIA is periodically trained on the users' devices and (2) a second phase where the latest joint model is used along with utility metrics to choose the best LPPM (among a set of LPPMs and corresponding configurations) each time a user wants to send a geo-located trace to the MCS. To avoid centralizing raw data in a trusted proxy server, we design FURIA using the FL paradigm [159].

We extensively evaluate EDEN using three real-world mobility datasets. We compare the performance of EDEN, both in terms of privacy and utility to the one of three off-the-shelf LPPMs using three configurations for each LPPM to cover the spectrum from strong privacy guarantees (despite the resulting impact on data utility) to strong utility objectives (with weaker privacy guarantees). For measuring the privacy offered by EDEN, we considered three state-of-the-art re-identification attacks (*i.e.*, POI-attack [193], PIT-attack [87] and AP-attack [152]) and combine them in a single, stronger attack that relies on majority voting. This attack is run on the MCS (considered as an adversary) and is different from FURIA. For measuring utility, we use two types of metrics: a quantitative metric and a qualitative metric. The quantitative metric, *i.e.*, area coverage, evaluates how far the area covered by a protected mobility data overlaps with the one of the original data. The qualitative metric captures the degradation in pollution measurements taken from a fourth real-world air pollution dataset [31]. An additional qualitative metric is considered, *i.e.*, range queries. It counts the number of users going through regions. This metric is useful for analyzing traffic congestion in a city, and it is taken from a fifth cab drivers mobility dataset in the city of San Francisco. In addition to comparing EDEN to state-of-the-art LPPMs, we consider two extreme solutions: a Privacy Oracle, which knows the attack run by the MCS and chooses the best LPPM accordingly and a Utility Oracle (referred to as NOBF), which sends pseudonymized raw data to the

MCS without applying any LPPM. The results show that EDEN provides a better tradeoff between privacy and data utility compared to individual LPPMs with a fair behavior for almost the users in the system (*i.e.*, users with similar mobility patterns receive similar privacy gain by EDEN).

The work proposed in this chapter has been published and presented in Ubicomp Conference 2021 [128].

**Roadmap** The remainder of this chapter is structured as follows. First, we illustrate our research problem in Section 5.2. Then, we describe the system model and an overview of our solution in Section 5.3. Further, in Section 5.4, we present a detailed description of EDEN and FURIA. An experimental evaluation of our solution is then presented in Section 5.5 and finally, we discuss the privacy limitations of EDEN in Section 5.6 and conclude the chapter in Section 5.7.

## 5.2 Problem Illustration

In crowd sensing applications, users contribute mobility data, which contains the user ID (*e.g.*, the device MAC address), the user location (*e.g.*, GPS latitude and longitude), the time at which the data has been collected, and the actual environmental measurement (*e.g.*, $NO_2$ measurements). Despite the pseudonymization of the user identity and techniques to hide the IP address of the originating device (*e.g.*, by using anonymous communication protocols such as TOR [205]), sharing mobility data may still leak information about users as discussed in the following section.

Consider an app developer, say Bob, who has to integrate privacy-by-design in a crowd sensing application. Bob needs to choose an LPPM with the appropriate configuration to protect users' mobility data before sharing it with a MCS. Bob does not want to implement yet another LPPM as there already exists a variety of LPPMs proposed by the research community. However, the actual level of privacy offered by each LPPM and the impact of the latter on data utility can dramatically vary according to how these LPPMs are configured.

To better illustrate this problem, we perform (on behalf of Bob) an experiment

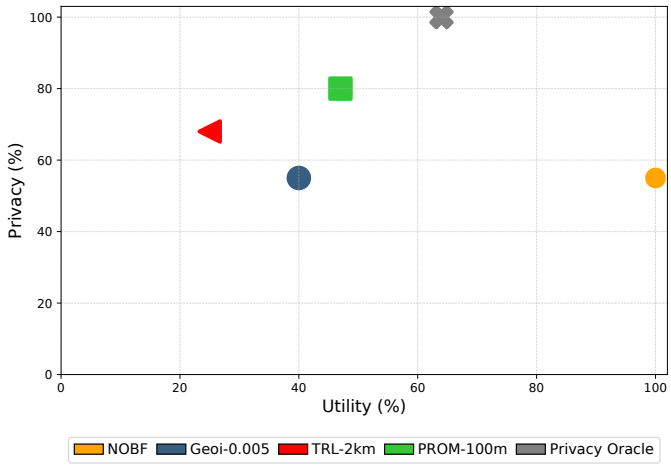


**Figure 5.1:** *Impact of LPPMs on privacy vs. utility tradeoff in the Privamov dataset.*

on the Privamov dataset [169]. We compare the privacy *vs.* utility tradeoff offered by three state-of-the-art LPPMs: Geo-Indistinguishability [21], Trilateration [119] and Promesse [196] noted *Geoi, TRL* and *PROM*, respectively. In order to find a satisfactory compromise between privacy and utility, we configure each of these LPPMs with an average privacy level, *i.e.*, $\epsilon = 0.005$, $r$=2km and $\alpha$=100m for *Geoi, TRL* and *PROM*, respectively. We provide more details about further configurations in Section 5.5.1. In addition to these three LPPMs, we also evaluate two alternatives: NOBF, which represents pseudonymized raw data without any additional obfuscation and Privacy Oracle, which is an oracle that selects the best LPPM for each individual trace (the LPPM which prevents re-identification and maximizes the data utility).

We measure privacy as the ratio of user's mobility data which is not re-identified by the MCS-side attacker over all mobility data. The attacker applies a majority voting using three state-of-the-art attacks, namely, POI-Attack, PIT-Attack and AP-Attack. We refer to this attack as *Mv-Attack* (for majority voting attack). On the other hand, we evaluate utility with the area coverage (AC) metric [198]. In Figure 5.1, an aggregate value of AC is depicted in the x-axis. It is computed as in Equation 5.1; where $F = \begin{pmatrix} 0 & 0.25 & 0.5 & 0.75 & 1 \end{pmatrix}$ refers to the vector of utility factors and $U = \begin{pmatrix} u_0 & u_{0.25} & u_{0.5} & u_{0.75} & u_1 \end{pmatrix}$ refers to the data proportion with $AC = 0$ or $AC \in ]0\ 0.25]$ , $]0.25\ 0.5]$ , $]0.5\ 0.75]$ and $]0.75\ 1]$ of the raw mobility data, respectively.

$$Utility = F^T.U \tag{5.1}$$

Note that Equation 5.2 is verified:

$$\sum_{i \in F} u_i = 1 \tag{5.2}$$

If no LPPM protects against *Mv-Attack*, the Privacy Oracle chooses to drop the data instead of sending it to the MCS. Privacy Oracle constitutes the best choice that can be done (from a privacy perspective) if Mv-attack is known by the defender and if all the data is centralized in a trusted proxy server. On the other side of the spectrum, NOBF is the best choice that can be done to preserve the data utility.

Results are depicted in Figure 5.1. In these results, we can observe that the Privacy Oracle protects 100% of the data and provides 65% of AC, which is better both in terms of privacy and data utility compared to individual LPPMs. In practice, the Privacy Oracle finds an LPPM resisting the attack for 96% of the traces and drops 4% of the remaining traces.

Therefore, we conclude that a Privacy Oracle has the potential to outperform all other state-of-the-art LPPMs in terms of privacy *vs.* utility tradeoff. However, the latter assesses the privacy using the *Mv-Attack*, which needs to centralize the raw data to a proxy server. The goal of this chapter is to design a solution that is as close as possible to the Privacy Oracle without centralizing the raw data in a proxy server.

## 5.3 EDEN Design Principles

In the following, we first present an overview of EDEN and its architecture in Section 5.3.1 and then we describe the threat model we considered in Section 5.3.2.

### 5.3.1 EDEN Overview

EDEN is a user-side mobility data protection system for crowd sensing applications. Its architecture is depicted in Figure 5.2. To better illustrate how it operates, let us consider Alice (depicted in the center of the figure), a participant in a crowd sensing campaign. Along her journey, Alice collects geo-located environmental measurements.

At a given point in time, the crowd sensing application decides to send the collected data (depicted at the bottom of the figure) to the MCS. Before sending this data to the MCS, EDEN automatically sanitizes the data without the implication of Alice. It applies a given LPPM among a set of available choices. For each LPPM, EDEN considers various configurations going from configurations that enforce strong privacy to ones that rather try to preserve data utility. Specifically, EDEN applies each LPPM to the raw trace and evaluates both: (1) the re-identification risk of the trace using this LPPM and (2) the corresponding data utility. For evaluating the re-identification risk, EDEN uses FURIA, a federated learning model, trained as depicted in the left part of Figure 5.2. When Alice's device fulfills a set of predefined requirements to participate in the FURIA training process (*e.g.*, her device is idle, charging, and connected to WiFi), it downloads from a server called the FURIA Master Server, the latest FURIA global model, trains the model using its own collected data and sends the updated model back to the server. Once the server receives a predefined number of users' responses, it aggregates by averaging them and provides a new version of FURIA model. This way, FURIA continuously learns and dynamically improves its global knowledge with new discriminating mobility patterns of incoming users. The training process of FURIA is performed by night independently from the process of protecting mobility traces using EDEN, which is illustrated on the right part of Figure 5.2. In this part of the figure, EDEN prepares batches of protected mobility data and sends a batch periodically (if any) to the MCS. This batch of geo-located data has been protected by using an LPPM for which the re-identification risk assessment performed using FURIA passed (*i.e.*, FURIA could not re-identify Alice as the originating user of this trace). If two LPPMs (or two variants of the same LPPM) pass the risk assessment, the one that has the best data utility is chosen. If no LPPM passes the FURIA risk assessment, then EDEN makes a decision according to a given configuration policy. For instance, it would drop the data if it is configured with a conservative policy. We describe other policies than the conservative one in Section 5.4.2. The configuration policy choice in EDEN can be set by the participant according to their preferences.

In this chapter, we considered three state-of-the-art LPPMs, each having three configurations, and we used three utility metrics further described in Section 5.5. Though, EDEN is not tight to a given set of LPPMs or utility metrics. More LPPMs, with their corresponding configurations and more utility metrics whether quantitative or application-dependent can be easily used in EDEN.

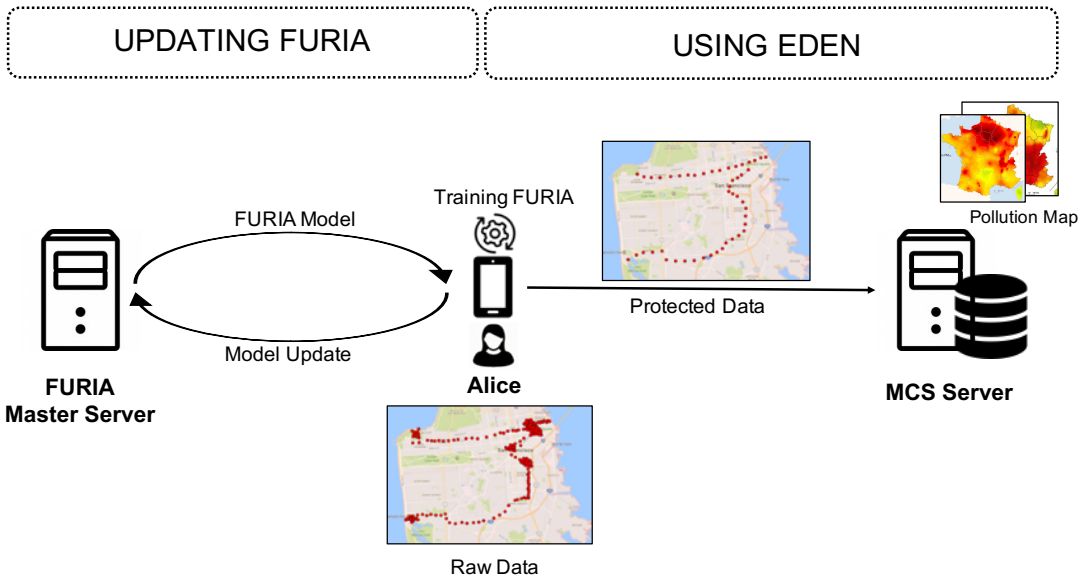


**Figure 5.2:** EDEN *overview*

### 5.3.2 Threat Model

As depicted in Figure 5.2, EDEN uses FURIA to assess the re-identification risk of a protected trace. FURIA is designed following a traditional client-server FL protocol using a master server. The communication channels between the clients and this server are encrypted.

Furthermore, we assume that users' devices are trusted and that the data protected by EDEN is sent anonymously to the MCS (*e.g.*, using an anonymous communication protocol such as TOR [205]). We consider the MCS as an adversary and we assume it to be *honest-but-curious* (*i.e.*, semi-honest [191]). Specifically, the MCS collects and processes geo-located environmental measurements to produce aggregate data for its clients. It performs this task honestly, *i.e.*, without deleting or altering the received data. However, it is curious because he may exploit the received data to learn valuable private information which may interest him or any third party (to whom he might sell the data for advertising purposes). For instance, the adversary may conduct inference attacks on the received data and consequently reveal the user identity or other sensitive attributes (*e.g.*, POIs, social links, *etc.*,), even if the participant is using the application anonymously [131]. Specifically, we assume that the MCS tries to link the received data to known mobility profiles he has previously built based on

leaked mobility data. To reach this objective, we consider that the MCS implements the latest available user re-identification attacks he found in the literature. Precisely, it combines three state-of-the-art attacks, namely, AP-Attack [152], POI-Attack [193] and PIT-Attack [87] in a single, stronger majority voting attack (*Mv-Attack*). Then upon receiving a trace protected by EDEN, *Mv-Attack* runs the three attacks and performs a majority voting between the predicted values of these attacks and returns the identity label that received more votes.

## 5.4 EDEN Detailed Description

In this section, we dive into a detailed description of how the re-identification risk assessment is done with FURIA in Section 5.4.1 and how mobility traces are protected with EDEN in Section 5.4.2.

### 5.4.1 Re-identification Risk Assessment with FURIA

FURIA's global model is a crucial part of our contribution. It applies FL principles to build a re-identification risk assessment model. The latter learns discriminating mobility patterns that uniquely identify users and helps our system to assess LPPMs in a privacy-preserving way. As depicted in Figure 5.4, FURIA involves two parties: (i) mobile user devices where raw mobility data is stored and where model updates are computed, (ii) the FURIA Master Server where model updates provided by various users are aggregated. FURIA operates as follows. First, the FURIA Master Server initializes a classification algorithm with random parameters. In this chapter, we use multi class logistic regression, a simple yet effective multi class classification method that satisfies very well our objectives, after an empirical experiment. Precisely, we compared three methods of classification, namely, multi class logistic regression (LR) [160], random forest (RF) [45] and a multi-layer perceptron neural network (NN) [180]. Figure 5.3a shows the re-identification rate over three unprotected datasets: Geolife, MDC and Privamov, described in Section 5.5.1. LR is slightly better than RF with +3% of re-identification rate in Geolife and both LR and RF are better than NN with up to +11%. Thus, the retained model for the rest of our work is LR.

The FURIA Master Server sends this model to all participants (step ① and ② in Figure 5.4). This model is denoted as $AF_0$. Each participant $U_j$ transforms its raw mobility data of the day to feature vectors (step ③) and trains the model $AF_0$ locally on the generated feature vectors. Once all participants have finished the first learning round, they send their local updates (*i.e.*, gradients) to the FURIA Master Server, (step ④). Upon receiving model updates, the FURIA Master Server aggregates users' gradients and produces an updated model, denoted $AF_1$, ready to use at the following day. This process is iteratively done and generally takes place at night time in order to avoid any interference with other applications running on the user's device. Indeed, model updates are computed when the user's device is idle, plugged in and connected to WiFi, which is generally the case at night time. FURIA processing is inspired by active/online learning where daily mobility data is unrolling between the train set of the current learning round and the testing set of the following one.

In FURIA, three types of features are considered: (i) spatial features (ii) temporal features and (iii) aggregated features. ***Spatial features***. To synthetically capture spatial information, records (*i.e.*, *lat* and *lng*) are projected on a heat map. A heat map is a set of cells of equal size. For each cell, the proportion of mobility records in a given trace $T$ that belongs to that cell is computed. This corresponds to the cell visit rate. ***Temporal features.*** The temporal information is considered to differentiate similar mobility patterns but at different times of the day. In FURIA we use the average time of the day of all the records in $T$. This is convenient in the case of crowd sensing applications, where mobility traces are generally in the order of minutes or hours length without exceeding a day. ***Aggregated features.*** Other types of information are extracted to enrich the user mobility profile. For instance, the number of mobility records in $T$ is considered. This allows to represent service usage intensity. We also extract the centroid of the mobility trace $T$ (*i.e.*, centroid's latitude and longitude), to capture a central position of the user's mobility in a map.

The spatial features are usually considered in related work [152, 138, 172]. In addition to this type of features, we explore several other types of features (*e.g.*, temporal and aggregated features) and evaluate their benefits for the user re-identification model, as presented in Figure 5.3b. The results show that temporal and aggregated features improve the re-identification rate of the attack. Specifically, the use of combined spatial, temporal and aggregated features increases the re-identification

success rate with +2%, +7% and +8% in comparison to the use of only spatial features in Privamov, Geolife and MDC datasets, respectively.

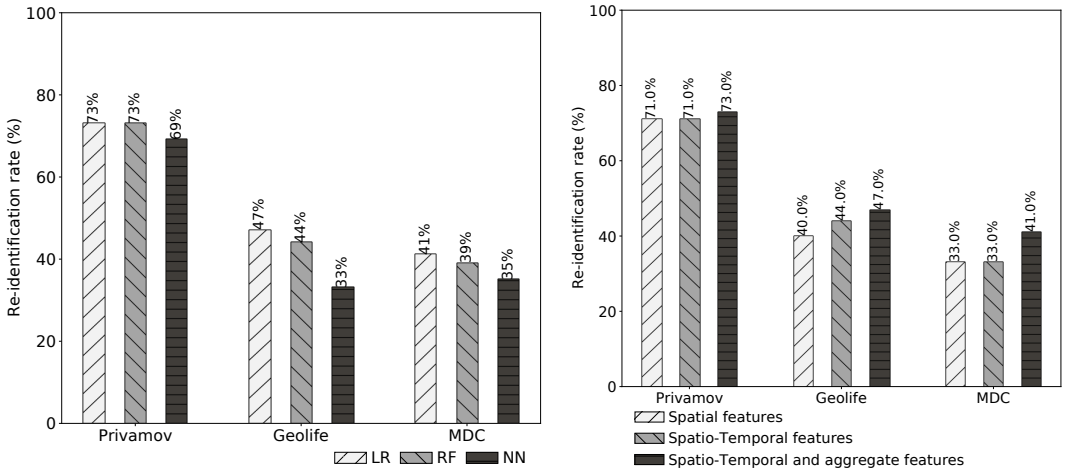


**(a)** *Model variation* **(b)** *Feature variation*

**Figure 5.3:** *Empirical experiments.*

### 5.4.2 Protecting Mobility Traces with EDEN

The detailed architecture of EDEN is depicted in Figure 5.5 and its behavior is described in Algorithm 3. EDEN takes as input a user mobility trace $T$ and a set of LPPMs $\mathbb{L}$ with various configurations (*i.e.*, low, medium, and high impact on privacy *vs.* utility tradeoff) and returns as output a protected mobility trace $T'_i$, which will be sent to the MCS. EDEN has four main components, the first component *Apply LPPMs* applies exhaustively all the LPPMs implemented in EDEN with their configuration variants on the raw mobility trace $T$ stored on the user smartphone (step ① in Figure 5.5). As a result, it produces a set of obfuscated versions of the same raw mobility trace *i.e.*, $C = \{T'_1, T'_2, \ldots, T'_n\}$. The second component *Format Data* transforms the different obfuscated traces available in the set $C$ into feature vectors (step ② in the figure). The third component *Global Model* FURIA uses the latest version of the model $AF_i$ as a privacy risk assessment metric. Specifically, it evaluates the user re-identification risk of each feature vector of the obfuscated data, (step ③). If the model fails in predicting the right identity label associated with the transformed mobility trace, the latter is potentially elected to be sent to the MCS.

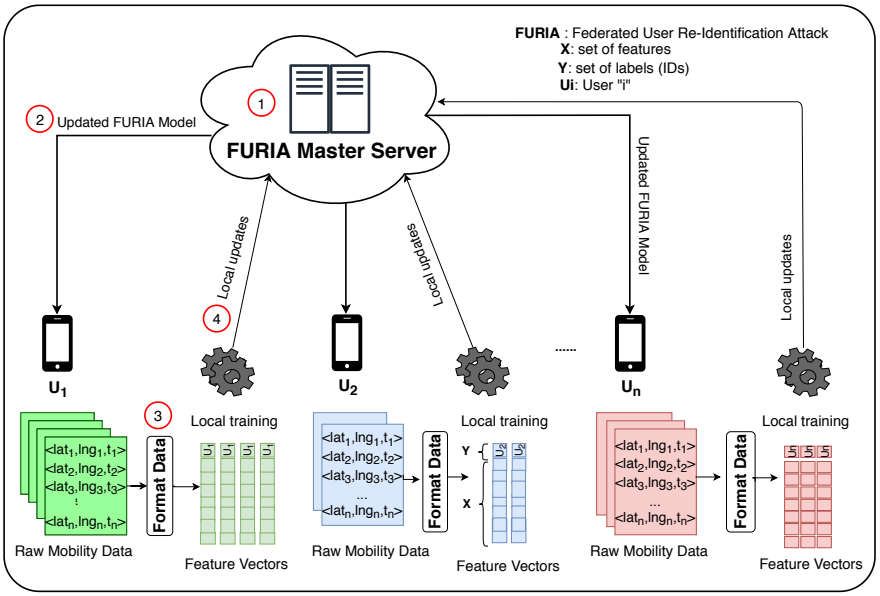


**Figure 5.4:** FURIA *architecture*

Otherwise, if the model succeeds in predicting the right identity label associated with the mobility trace for all the considered LPPMs, three different policies can be adopted by EDEN (step ④) as described later.

Finally, the last component in EDEN is *Best Coverage*. It selects the protected mobility trace candidate that maximizes data utility. EDEN can consider various utility metrics. In this chapter, we use the area coverage metric (AC) [198]. It is computed between the original and the obfuscated mobility trace, $T$ and $T'_i$ respectively, to measure how much the alteration caused by an LPPM affects the regions visited by a user (step ⑤). We provide more details about the AC metric in Section 5.5.1. Finally, the obfuscated mobility trace that better resists the re-identification test performed by FURIA and that has the best utility is sent to the MCS. The latter processes the received data and produces useful information for users (step ⑥).

### EDEN Policies

Three policies are considered by EDEN if a mobility trace is re-identified by FURIA. The first policy is **EDEN-pessimistic** (EDEN-pes): it is the most conservative policy as it simply deletes a mobility trace that FURIA is able to re-identify. The

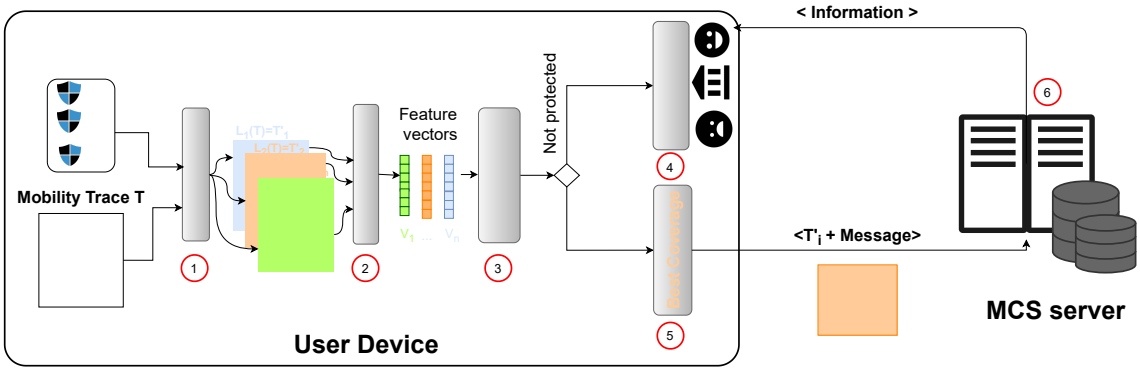


**Figure 5.5:** EDEN *architecture*

rationale behind this policy is that if FURIA is able to re-identify a mobility trace, an external attacker could very well reach the same result. The downside of this solution is that it causes data loss from the application perspective.

The second policy is **EDEN-optimistic** (EDEN-opt): an opposite solution to the previous one where the mobility traces that FURIA is able to re-identify are also sent to the MCS without applying any LPPM or by applying a default LPPM. We use this policy as a baseline to assess the impact of sending traces despite the red flag raised by FURIA regarding the re-identification risk of some mobility traces.

The third and last intermediary policy is **EDEN-balanced** where sending or dropping mobility traces is based on a local metric evaluating how a considered mobility trace is similar to past mobility traces of the same user. Towards this purpose, we use the Topsoe divergence metric [52]. The latter is computed between two probability distributions : (1) the heat map of the current raw mobility trace $T$ and the heat map corresponding to the past mobility data of the same user, which is stored on the user's device, (see Equation 5.3). It is a derived symmetric version of the Kullback Leibler divergence [52] which measures the information deviation between the user's past mobility and the current one. If the deviation is high (greater than a threshold empirically set to 0.8), the mobility trace is sent to the MCS, otherwise, it will be deleted. The set value of the threshold fits well the distribution of Topsoe deviation values.

$$d_{Topsoe}(P,Q) = \sum_i \left[ P_i \ln\left(\frac{2P_i}{P_i+Q_i}\right) + Q_i \ln\left(\frac{2Q_i}{P_i+Q_i}\right) \right] \tag{5.3}$$

**Algorithm 3** EDEN algorithm.

**INPUT:**

$T$: mobility trace, $\mathbb{L}$: a set of LPPMs, $AF_i$: $i^{th}$ model, $AC$: utility metric, *policy*: {*EDEN-opt, EDEN-pes, EDEN-balanced*}, $H$: past user data, $\delta$: deviation threshold

**OUTPUT:**

$T'$: protected mobility trace.

**function** $\mathcal{EDEN}(T,\mathbb{L},AF_i,AC,$ *policy*, $H,\delta)$
  $Candidates \leftarrow \emptyset$
  **for** $\mathcal{L}$ **in** $\mathbb{L}$ **do** Apply LPPMs
    $T' \leftarrow \mathcal{L}(T)$
    $V' \leftarrow FormatData(T')$ Transform T' to a Feature vector V'
    **if** $AF_i(V') \neq U$ **then** $Candidates \leftarrow Candidates \cup \{T'\}$ Re-identification Risk assessment
  **end for**
  **if** $Candidates \neq \emptyset$ **then**
    **return** $\{\underset{T' \in Candidates}{\arg\max} (AC(T,T'))[0]\}$
  **else** EDEN Policies
    **if** *policy* ="EDEN-pes" **then return** $\emptyset$
    **if** *policy* ="EDEN-opt" **then return** $T$ or $\mathcal{L}(T)$.
    **if** *policy* ="EDEN-balanced" **then**
      $P \quad Heatmap(T)$
      $Q \leftarrow Heatmap(H)$
      $deviation \quad d_{Topsoe}(P,Q)$
      **if** $deviation > \delta$ **then**
        **return** $T$
      **else**
        **return** $\emptyset$
      **end if**
    **end if**
  **end if**
**end function**

## 5.5 Experimental Evaluation

We start this section by describing our implementation and experimental setup in Section 5.5.1 and then, our evaluation answers the following questions:

- How does EDEN perform against an adversary attack compared to competitors? (Section 5.5.2)
- What is the impact of EDEN on data utility compared to competitors? (Section 5.5.3)
- What are the LPPMs used by EDEN to effectively protect mobility data? (Section 5.5.4)
- What is the run-time overhead of EDEN? (Section 5.5.5)
- How EDEN is fair in the protection of mobility data of similar users? (Section 5.5.6)

### 5.5.1 Implementation and Experimental Environment

#### Experimental Setup

All the experiments related to the MCS attacker are carried out on a server running an Ubuntu 14.04 OS with 50GB of RAM and 16 cores of 1.2GHz each. Both EDEN and FURIA are developed in Python using the Pytorch library [10]. We used S2Geometry library [11] for the decomposition of the map into cells of approximately equal size. The cell edge length ranges from 212m to 296m. To accelerate the training process of our federated learning model, we use a machine with NVIDIA TESLA V100 GPU. Participating users are simulated by considering the data coming from real-world mobility datasets.

### Mobility Datasets

In our experiments, we use three real-world publicly available mobility datasets with a summary given in Table 5.1. These datasets are: Geolife [257], MDC [135] and Privamov [169]. In our experiments, we extracted only the most active month (*i.e.*, 30 days) of each dataset for a fair comparison.

In the context of location privacy, these datasets are used by many state-of-the-art LPPMs in order to assess the effectiveness of their approach [196, 127, 153, 193]. That is why, we decide to evaluate our approach on these datasets to be in line with the research community.

**Table 5.1:** *Description of datasets*

| Name | Geolife | MDC | Privamov |
|---|---|---|---|
| # users | 42 | 144 | 48 |
| location | Beijing | Geneva | Lyon |
| # records | 1,468,989 | 904,282 | 774,401 |
| area ($km^2$) | 16,808 | 41.37 | 47.87 |

### Evaluation Scenario

We simulate mobile users that correspond to the users of the previously described datasets. We assume that the data corresponding to the first 15 days of each dataset has been leaked to the MCS. Using this data, the MCS builds user profiles. The same data is used to train the first FURIA model (*i.e.*, $AF_0$) to be in the same conditions as the adversary. The remaining 15 days of each dataset are then used as a test set. Specifically, FURIA is inspired by active/online machine learning. Its training/testing mobility data is unrolled with a time window of 24 hours. For example, the $16^{th}$ day mobility data is used to test $AF_0$ and to train $AF_1$, the $17^{th}$ day mobility is used to test $AF_1$ and to train $AF_2$ and so on. It means that our training set of mobility data is incremented day by day and many phases of test occur on each newly trained model $AF_i$.

Moreover, as sharing data with the MCS in real-time is energy-consuming [89, 236], we assume that the user's crowd sensing application prepares batches of 30 minutes

in length to be as close as real-time data transmission use cases, protects this data using EDEN and periodically sends it to the MCS.

Upon receiving a geo-located trace, the MCS uses this trace to update its target map. Simultaneously, the MCS tries to re-associate the received trace to one of the user mobility profiles it has previously built, using *Mv-Attack*. We compare EDEN to LPPMs that are either applied blindly on users' mobility data or assisted with a Privacy Oracle.

### Utility Metrics

To evaluate the impact of EDEN and its competitors on the quality of the generated data, there are two categories of utility metrics, proposed in the literature [199]. *(i) Data-centric or quantitative* metrics which measure the distortion between the original and the obfuscated mobility data. In this chapter, as previously mentioned in Section 5.4.2, we use the AC metric. To recall, it computes the overlap between the obfuscated and the original mobility trace using the F1-score. This metric is able to capture the degradation in data utility caused both by LPPMs that remove data points (*e.g.*, PROM), and the degradation caused by LPPMs that add data points or move them spatially (*e.g.*, TRL and Geoi). Thus, this metric can be used in various applications, such as in transportation mobile applications where a data analyst can use the AC metric to adapt the availability of public transportation in areas according to visiting user density. *(ii) Application-centric or qualitative* utility metrics, which compare the result of a given application before and after applying an LPPM. In this chapter, we consider two real-world use cases. In the first use case, we visualize the air pollution degradation map before and after the application of EDEN and state-of-the-art LPPMs on a crowd sensing air pollution dataset [31]. This map allows to detect areas where the level of gaseous pollutants is high (*e.g.*, hotspots of $NO_2$ or $CO$). This dataset is described in Section 5.5.3. In the second use case, we use range queries metric, a classical operation which compares the number of unique users who go through areas during a time window before and after obfuscation [196]. An illustrative example is provided in Figure 5.6, where two range queries **Q1** and **Q2** of different radius are performed (the temporal dimension is not represented). Before the obfuscation process, **Q1** and **Q2** return 3 and 1 users, respectively, whereas after obfuscation, **Q1** and **Q2** return 1 and 3 users. Thus, the

utility is measured as the range query distortion defined in [15]. In our example, the distortion of **Q1** is $\frac{|3-1|}{3} = \frac{2}{3}$ and the distortion of **Q2** is $\frac{|1-3|}{1} = 2$. Then, the average query distortion is computed, *i.e.*, $\frac{\frac{2}{3}+2}{2} \approx 1.333$.

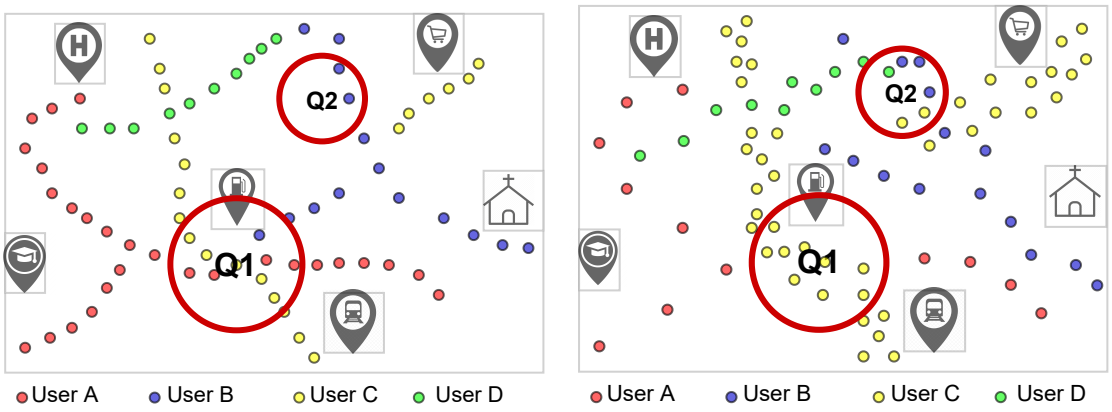


**(a)** *Before obfuscation* **(b)** *After obfuscation*

**Figure 5.6:** *Illustrative example of the range queries metric.*

### FURIA Configuration

To build FURIA, we opt for the multi class logistic regression (LR) algorithm. The training of our global model is done over multiple rounds ($R_i$). Each round represents a 1-day training, except the first round ($R_0$) where a training set of 15 days is used. The latter is considered as historical data stored on the user device and previously leaked to the adversary (*i.e.*, *Mv-Attack* has access to 50% of the mobility dataset). Thus, we decide to start the training process with the same knowledge of the adversary to be in the same conditions. We assume that users train the model at the end of the day (*e.g.*, at night time) with the data collected during that day in order to prepare the model of the following day. Thus, the model is actively trained by incoming mobility data and improves its global view day by day. In each round, we run 100 epochs/user with a variable batch size. We tune the batch size according to the number of collected traces per user participating in the given round. In addition, we use the stochastic gradient descent optimizer (SGD) with a learning rate of 0.001. These parameters are fixed empirically after several experiments.

### User Re-identification Attack Configuration

*Mv-Attack* is made by the combination of three state-of-the-art attacks, namely, AP-Attack [152], POI-Attack [193] and PIT-Attack [87]. By combining these attacks, we obtain an attack that is stronger than considering the attacks separately as the adversary gets more confidence about the result of the re-identification process. Each attack has a set of parameters, described below. POI-Attack and PIT-Attack have two parameters for the extraction of POIs [259]. These parameters are the diameter of the clustering area set to 500 meters and the minimum time spent inside a POI set to 5 minutes to accommodate small traces. AP-Attack has a configuration parameter that corresponds to the square cell size set to 800 meters (default value [152]).

### Competitors

To evaluate EDEN, we select three state-of-the-art LPPMs, namely, Geo-Indistinguishably (Geoi) [21], Trilateration (TRL) [119] and Promesse (PROM) [196]. We select these LPPMs because they can be run on the user side (*i.e.*, without any external knowledge about other users' mobility) and they provide diverse guarantees: differential privacy, dummy-based obfuscation, and POI erasure, respectively. Each LPPM has its own configuration parameters. These parameters have an impact on the privacy *vs.* utility tradeoff. In our experiments, for *Geoi*, we set the privacy parameter $\epsilon$ to 0.01, 0.005 and 0.001. A lower value of $\epsilon$ leads to a higher level of noise added to mobility records and consequently ensures a higher level of privacy. For *TRL*, there is a circular region with a radius of $r$, that surrounds the real location of the user. The chosen values of this parameter are 1 km, 2 km, and 3 km. A higher value of $r$ generates a bigger region for location dummies and consequently ensures a higher protection level. And finally, for *PROM*, which has a parameter $\alpha$ that specifies the distance between two successive mobility points, we set $\alpha$ to 50, 100, and 200 meters. A higher value of $\alpha$ leads to a larger distance between points in a mobility trace and thus ensures a higher protection level. However, the latter can cause serious data loss, especially if a mobility trace is recorded in a short distance over a short period. In addition to the above three LPPMs with their configurations, we also evaluate two baselines: a Utility-centric baseline, referred to as *NOBF*, which corresponds to sending the data to the MCS without obfuscation (*i.e.*, sending raw, pseudonymized

data); and a Privacy-centric baseline, referred to as *Privacy Oracle*. This baseline is only used in the evaluation of the privacy *vs.* utility tradeoff. It represents a solution where the selection of the best LPPM is driven by the attack performed by the adversary. As such, perfect privacy can be reached but the chosen LPPMs can still degrade data utility.

### 5.5.2 Evaluation of Data Privacy

In this section, we evaluate the effectiveness of EDEN in terms of privacy in comparison to state-of-the-art LPPMs. For that purpose, we measure the data protection rate of EDEN's variants and its competitors against *Mv-Attack*. To recall from Section 5.2, the data protection rate is the percentage of mobility data that is not re-identified by the MCS. Results are depicted in Figure 5.7.

From this figure, we observe that on the Privamov dataset (Figure 5.7a), 55% of the data sent without obfuscation (*NOBF*) is not re-identified by the MCS. This percentage is the same when *Geoi-0.01* is used and slightly increases when *Geoi-0.005* and *Geoi-0.001* are used (56% and 64% of protected data, respectively). This is due to the dependency between successive mobility records which makes the $\epsilon$-*Geoi* guarantee loses its power to $n * \epsilon$-*Geoi* ($n$ being the number of records). The most privacy-protective LPPMs from the literature are *TRL* with an increased range $r$ or *PROM* with a large distance $\alpha$ between points in the mobility trace. Specifically, the proportion of protected traces reaches up to 74% when *TRL-3km* is used and 91% when *PROM-200m* is used. In the case of EDEN, 85%, 86%, and 87% of protected mobility data are recorded with the optimistic, balanced, and pessimistic variants of EDEN, respectively. We notice that *PROM-200m* outperforms EDEN's variants with +5% on average, this is due to the fact that *PROM* is based on the re-sampling of mobility traces by erasing points according to the distance parameter $\alpha$. Thus, if a mobility trace does not exceed 50, 100, or 200 meters of traveling distance, the latter will be deleted. The deleted data is not sent to the MCS and is considered protected. However, this dramatically degrades data utility, as further discussed in Section 5.5.3.

In the Geolife dataset (Figure 5.7b), 61% of mobility data is naturally protected against the *Mv-attack*. The application of *Geoi* and *TRL* with their different

configurations, does not improve the protection rate compared to the baseline. Using *PROM-50m*, *PROM-100m* and *PROM-200m* increases the protection rate with +4%, +6% and +10%, respectively. This is also due to the suppression of chunks where the mobility data does not exceed 50, 100, or 200 meters in a lap of 30 minutes. In this dataset, applying EDEN significantly improves the data protection rate reaching up to 87% and 90% of protected mobility data with *EDEN-balanced* and *EDEN-pes*, respectively. Finally, in the MDC dataset (Figure 5.7c), 68% of the mobility data is naturally protected. The application of EDEN's variants improves the protection rate with +10% on average compared to the NOBF baseline. This result has the same trend for the other LPPMs except with *PROM* which provides a higher protection rate (+4% on average). This result is due to the suppression of mobility traces.

Here, we observe that *PROM-50m* provides better protection with the MDC dataset compared to Privamov and Geolife datasets. MDC involves 144 users in a quite small area (around 41 $km^2$). Such a high population density naturally reduces user re-identifiability which, thus, enables higher protection. In contrast, the Privamov dataset has three times fewer users than MDC, and the Geolife dataset has only 42 users in a large geographical area (16,808 $km^2$). This makes users in these two datasets more distinguishable and, thus, harder to protect.

### 5.5.3 Evaluation of Data Utility

As described in Section 5.5.1, we measure the data utility using two types of metrics: a quantitative metric which is AC metric and a qualitative metric where we capture the degradation in pollution measurements taken from a real-world air pollution dataset [31]. In addition, we measure the number of cab drivers going through regions taken from a real-world mobility dataset [190].

#### Utility Evaluation Using AC Metric

In Figure 5.8, we evaluate the data quality of sent/not sent data to the MCS (the privacy dimension is not considered in this figure). In this figure, the AC equal to 0 corresponds either to the data deleted by the LPPM (*e.g.*, the case of promesse) or

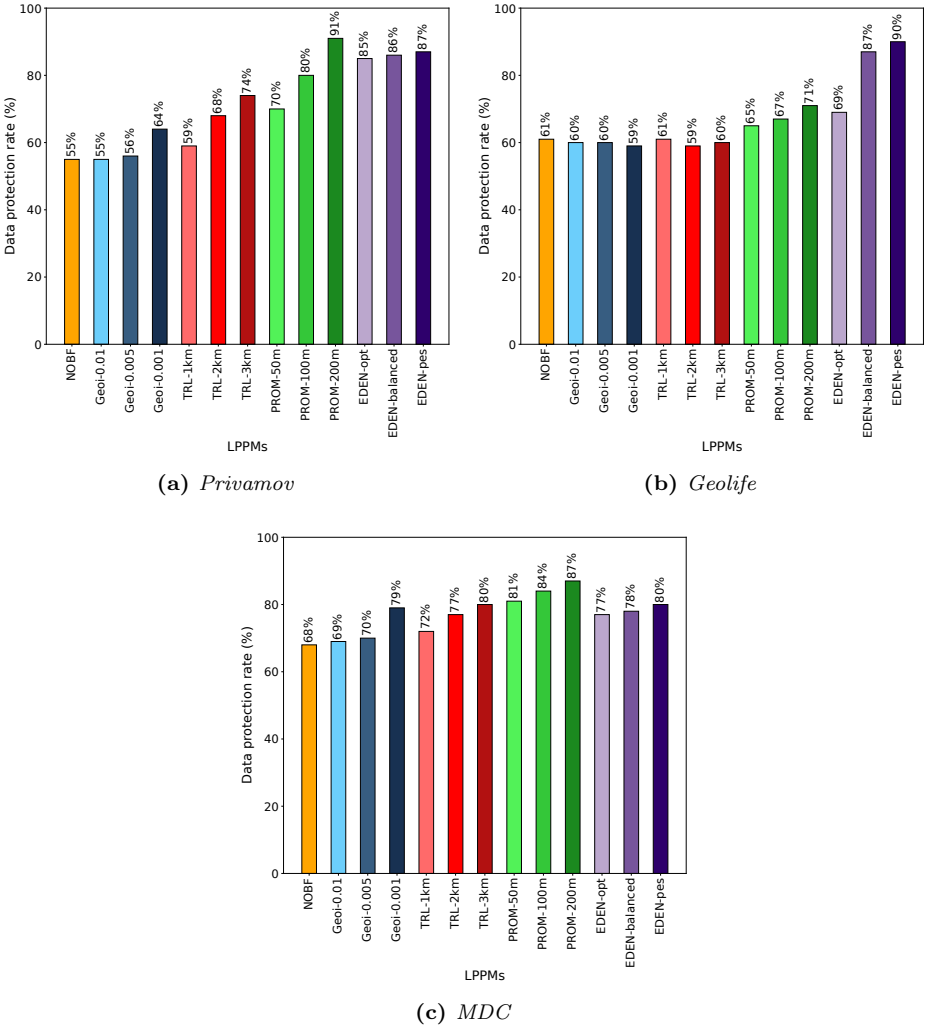


**(a)** *Privamov* **(b)** *Geolife*

**(c)** *MDC*

**Figure 5.7:** *Mv-Attack evaluation on EDEN vs. competitors.*

protected data that has no intersection with the original data. Further, we group AC values into four intervals. The best LPPMs are those that maximize AC in the interval $[0.75, 1]$ while minimizing the data loss ($AC = 0$).

In the Privamov dataset, we notice that EDEN's variants produce a balanced data quality. Specifically, EDEN-opt, EDEN-balanced, and EDEN-pes lose an average of 21.4% of mobility data and produce an average of 46% of data with an $AC < 0.75$. And finally, an average of 32% of the generated mobility data has an $AC > 0.75$. However, *PROM* with its different configurations has predominately darker bars. In particular, 50% and 74% of the generated mobility data has $AC$=0 with *PROM-100m*

and *PROM-200m* respectively. Promesse chooses to not share with the MCS a larger proportion of the generated data due to its suppression process. In contrast, the data chosen to be shared closely mirror the original mobility data: up to 65% of protected mobility data by *PROM-50m* has an $AC > 0.75$. *Geoi* and *TRL* degrade the quality of almost generated data. Specifically, in *Geoi-0.01*, *Geoi-0.005* and *Geoi-0.001*, an average of 92% of the resulting mobility data has an $AC < 0.75$. Even worst with *TRL* where an average of 95% of mobility data has an $AC < 0.5$.

In the Geolife dataset, the difference between EDEN's variants is more prominent. The data loss is reduced from 41% in EDEN-opt to around 10% in EDEN-pes. More than 79%, 58%, and 48% of the generated data by EDEN-opt, EDEN-balanced, and EDEN-pes, respectively, have an $AC > 0.75$. However, *PROM-50m*, *PROM-100m* and *PROM-200m* cause a data loss of 10%, 16% and 25%, respectively. The remaining data (*i.e.*, 85%, 76% and 62%) has an $AC > 0.75$. We observe that *PROM* has a better *AC* than EDEN. This is because EDEN prioritizes privacy over data utility. Finally, *Geoi* and *TRL* with their different configurations generate on average 89% and near 98% of mobility data with an $AC < 0.75$. Only 26% of the generated data by *Geoi-0.01* has an $AC > 0.75$.

In the MDC dataset, unlike *PROM* with its different configurations which cause 32%, 41%, and 50% of data loss, EDEN's policies reduce these amounts to 7%, 8% and 10%. They outperform all other LPPMs in terms of *AC* metric. Specifically, it protects an average of 67% of all data with $AC > 0.75$ compared to an average of 0%, 14% and 39% with *TRL*, *Geoi* and *PROM*, respectively.

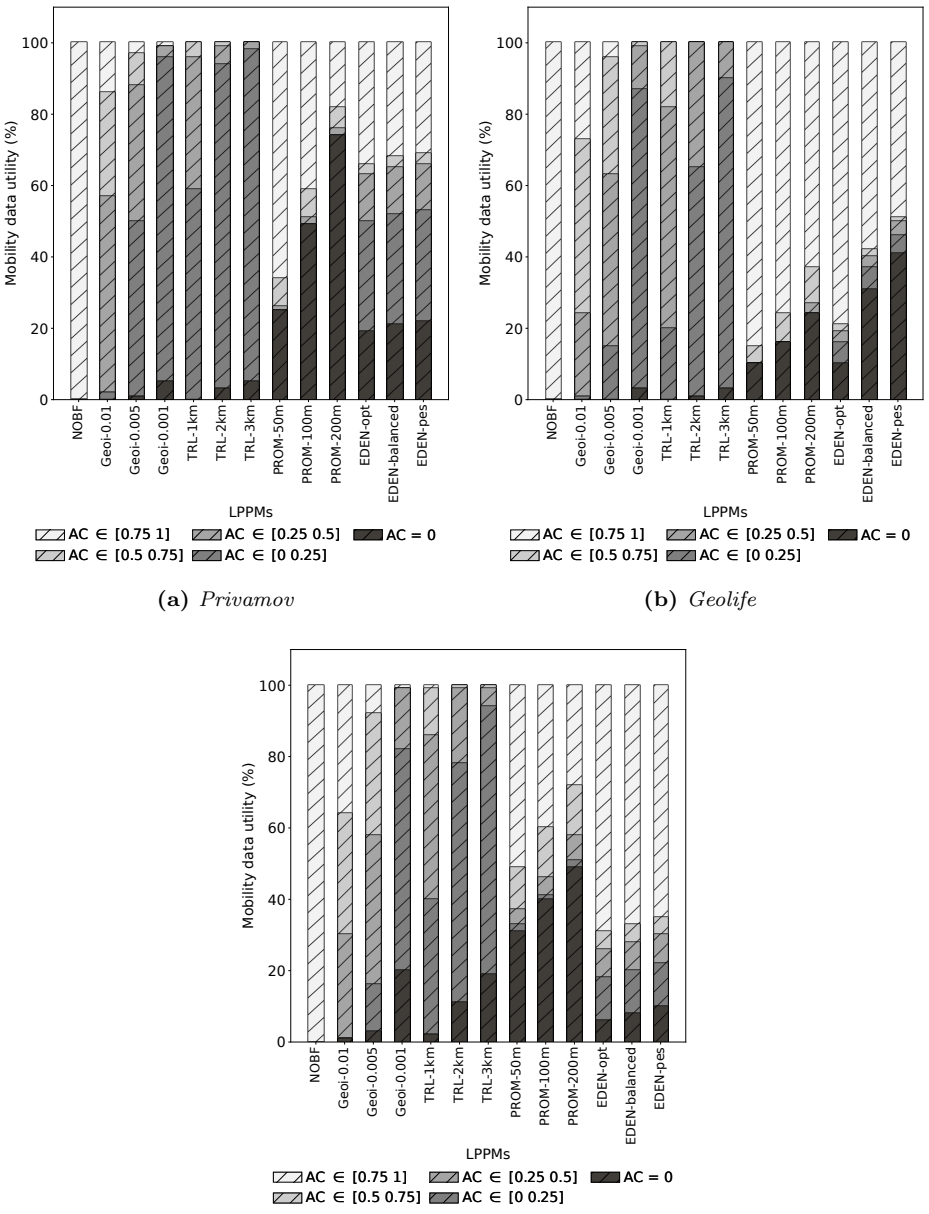


**(a)** *Privamov*

**(b)** *Geolife*

**(c)** *MDC*

**Figure 5.8:** *Impact of EDEN vs. competitors on the data utility using AC metric*

### Macro-benchmark on Air Pollution Measurements

We study the impact of LPPMs on the quality of air pollution data using a dedicated dataset [31]. In addition to mobility data, two application-specific measurements are

collected: the concentration of $NO_2$ and $CO$. The dataset involves 13 metropolitan bikes which have been equipped with pollution monitoring sensors for a duration of 112 days.

We compute the average $CO$ measurements over the duration when the $NO_2$ value is above $40\mu g/m^3$, which is the toxicity threshold defined by the WHO[1]. In Figure 5.9, we show the average $CO$ measurements when using the raw data (NOBF baseline) and after the application of EDEN and its competitors. We can observe that *Geoi* and *TRL* spread the measurements and, as the noise or the range increases, they create additional hotspots, *i.e.*, areas where the $CO$ value is high. The application of *PROM-50m* creates around six new hotspots where the level of $CO$ is now above the toxicity threshold. Thus, although the good privacy *vs.* utility tradeoff offered by *PROM-50m* on synthetic mobility data, including the $CO$ measurements yields poor performance. The use of *PROM-100m* can be harmful to public health because it eliminates existing high pollution hotspots from the original data. However, the application of EDEN can closely mirror the $CO$ measurements of NOBF; only one hotspot is missed.

**Macro-benchmark on the Number of Cab Drivers in San Francisco**

We study the impact of EDEN and its competitors on the mobility of 50 cab drivers from the Cabspotting dataset [190]. The objective, for example, is to find bottleneck locations that cab drivers go through in the city of San Francisco. For that purpose, we use the range queries metric, previously defined in Section 5.5.1. It counts how many unique users cross an area during a time window. We choose time windows ranging from 2 hours to 8 hours and circle areas whose radius range from 500 meters to 5,000 meters. We report the average query distortion in Table 5.2, which is the average distortion over 1,000 randomly generated queries. The results show that EDEN provides the smallest average distortion with 0.55% in comparison to, respectively, PROM which can reach 1.39%, TRL which can reach 16.45%, and Geoi which can reach 19.88%.

[1]World Health Organization, URL: https://www.who.int/fr/news-room/fact-sheets/detail/ambient-(outdoor)-air-quality-and-health

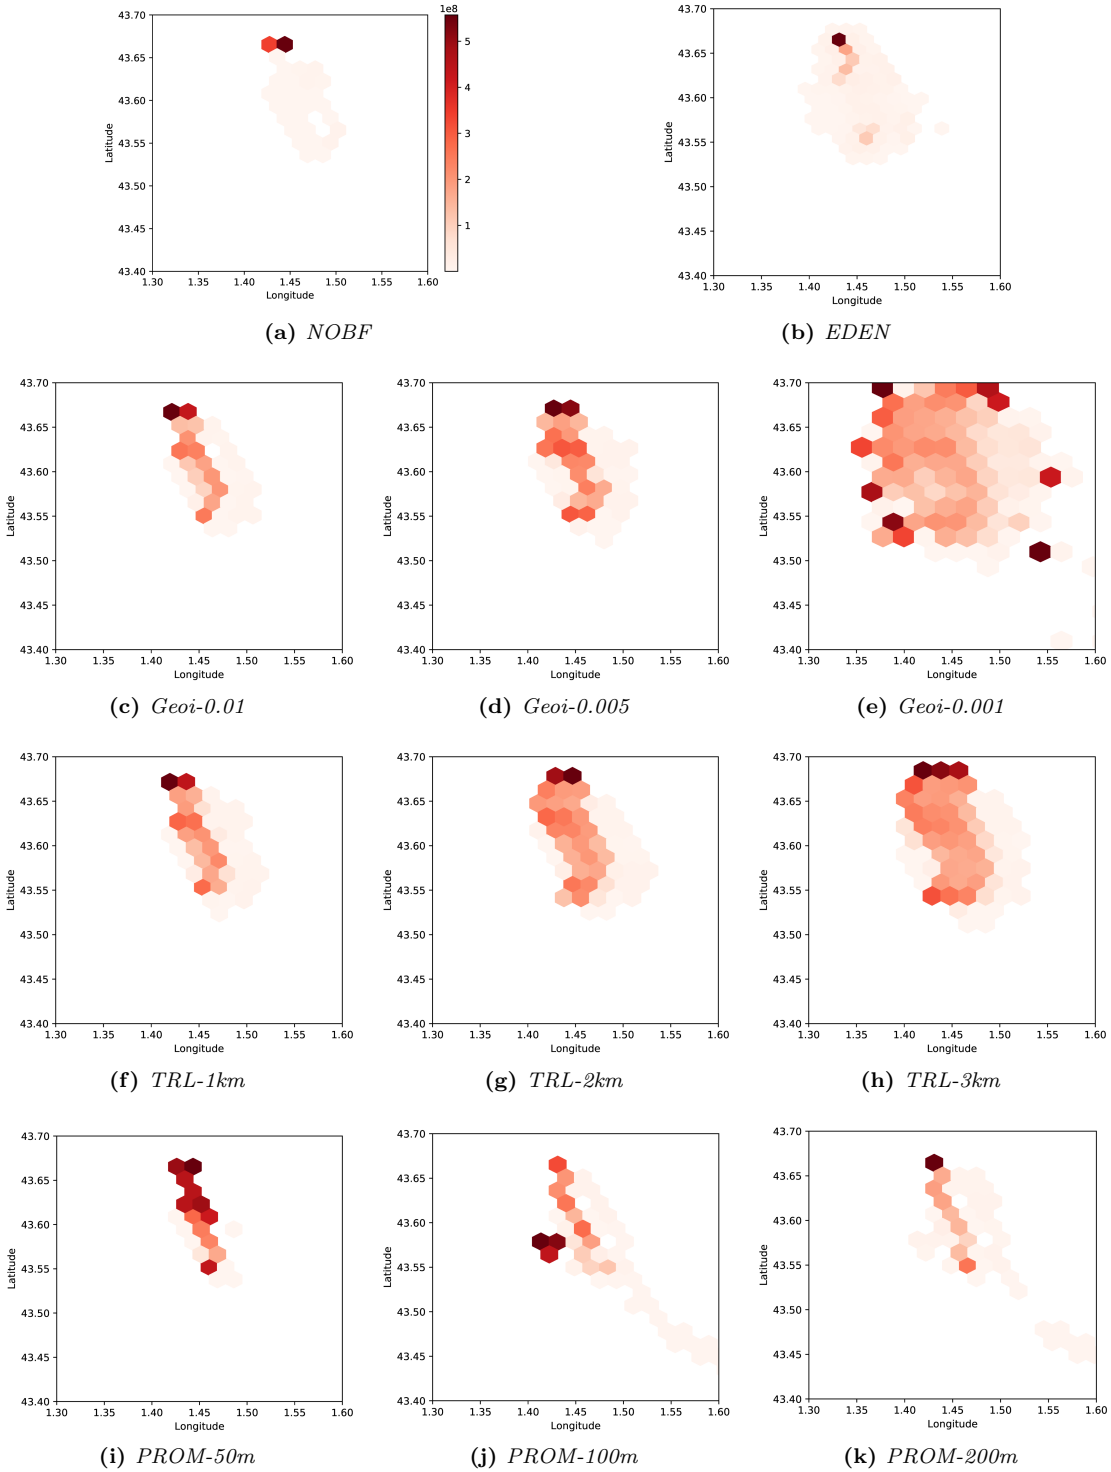


**(a)** *NOBF* **(b)** *EDEN*

**(c)** *Geoi-0.01* **(d)** *Geoi-0.005* **(e)** *Geoi-0.001*

**(f)** *TRL-1km* **(g)** *TRL-2km* **(h)** *TRL-3km*

**(i)** *PROM-50m* **(j)** *PROM-100m* **(k)** *PROM-200m*

**Figure 5.9:** *Macro-benchmark on air pollution dataset (CO gas)*

**Table 5.2:** *Average query distortion of* EDEN *vs. competitors.*

| **LPPM** | **Average Query Distortion** |
|---|---|
| EDEN | 0.55% |
| PROM-50 | 1.1% |
| PROM-100 | 1.28% |
| PROM-200 | 1.39% |
| TRL-1 | 6.96% |
| TRL-2 | 12.13% |
| TRL-3 | 16.45% |
| Geoi-01 | 1.07% |
| Geoi-005 | 2.73% |
| Geoi-001 | 19.88% |

### Privacy *vs.* Utility Tradeoff

In Figure 5.10, we evaluate the privacy *vs.* utility tradeoff of EDEN compared to its competitors. To this end, we use a scatter plot where a point corresponds to an LPPM configuration defined by two coordinates x and y; x represents an aggregate value of the AC utility metric computed as in Section 5.2; y represents the percentage of protected data (*i.e.*, data that is not re-identified by the MCS-side attacker). We consider that an LPPM reaches a good privacy *vs.* utility tradeoff if it belongs to the top right gray rectangle ***"R"***, *i.e.*, a rectangle where only LPPMs that have a protection rate greater than 80% with an AC greater than 40% are considered. We represented the *Privacy Oracle* in order to show the maximum utility that can be attained by a specific dataset when all mobility traces are protected. In the same vein, *NOBF* represents the maximum privacy that can be attained by a specific dataset when no LPPM is applied.

In the Privamov dataset, as depicted in Figure 5.10a, EDEN's variants (illustrated by purple stars) provide the best privacy *vs.* utility tradeoff. Specifically, All EDEN's variants belong to ***"R"*** whereas only *PROM-100* (illustrated by a green square) belongs to the low border of ***"R"***. In the Geolife dataset, as depicted in Figure 5.10b, only EDEN-pes and EDEN-balanced belong to ***"R"***. However, the rest of the LPPMs have a privacy value concentrated around 62% and a utility varying from 26% to 88%. Finally, in the MDC dataset, as depicted in Figure 5.10c, EDEN-pes and *PROM* with

its different configurations are inside ***"R"***. However, all EDEN's variants achieve a better privacy *vs.* utility tradeoff with a privacy level close to *PROM* but a utility value close to *Privacy Oracle*. To conclude, in the three considered datasets, EDEN achieves a better privacy *vs.* utility tradeoff than any other individual LPPM with at least one of EDEN's variants belonging to the target rectangle ***"R"***.

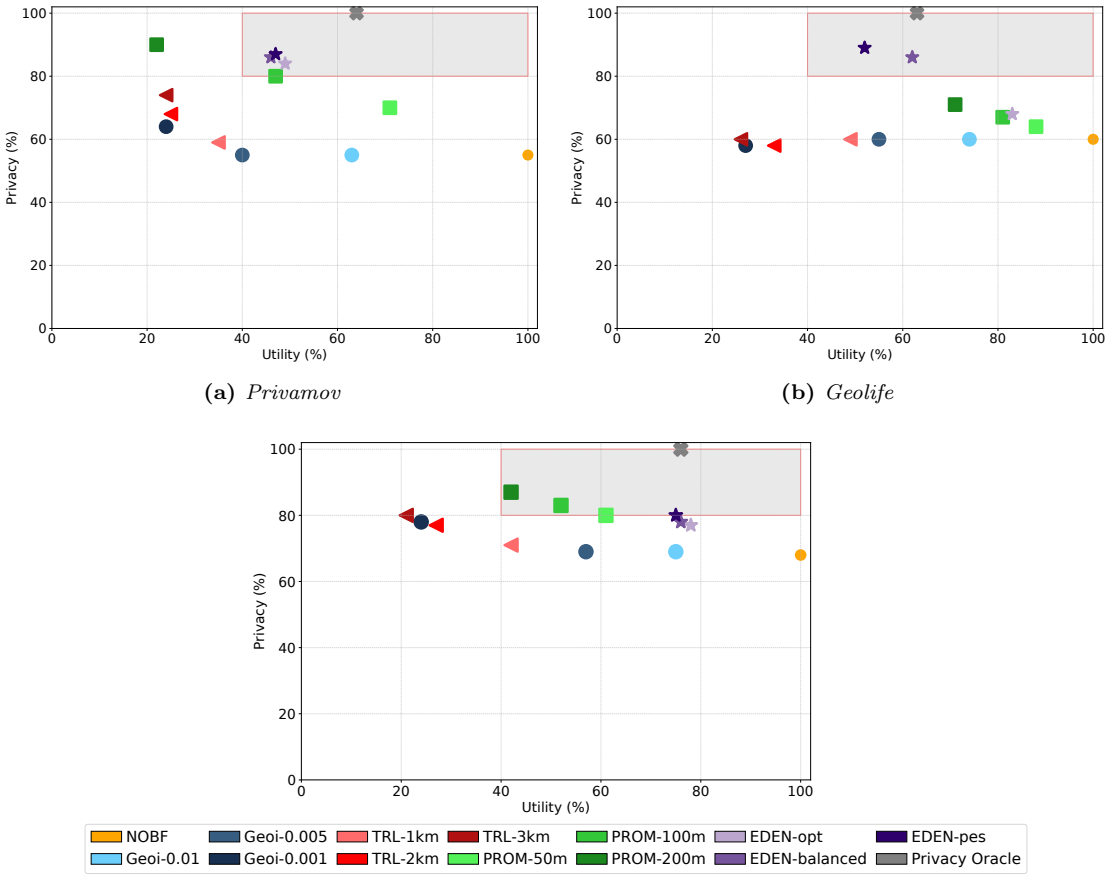


**(a)** *Privamov* **(b)** *Geolife*

**(c)** *MDC*

**Figure 5.10:** *Privacy vs. utility tradeoff.*

### 5.5.4 Fine-Grained Analysis of EDEN

In this section, we perform a fine-grained analysis of EDEN's variants. Unlike the other LPPMs, which are applied individually and blindly on the whole mobility data, EDEN protects users' mobility traces in a fine-grained way by choosing the most appropriate LPPM for each mobility trace. In Figure 5.11, we represent EDEN's

variants with multi-color bars where the hashed part is the re-identified portion of the mobility data, and the plain part is the portion of protected data against *Mv-Attack*. The color black represents the mobility data that FURIA is always able to re-identify regardless of the chosen LPPM. This data is deleted and considered as data loss from the MCS point of view. Each LPPM is represented with a different color and its intensity expresses the strength of the LPPM's configuration: darker colors are used for stronger LPPM configurations. From this figure, we observe that orange is the dominant color that refers to NOBF. Indeed, when FURIA is not able to identify a raw mobility trace, EDEN prioritizes this choice because it provides the maximum utility.

By focusing on the Geolife dataset, we observe that the data loss of *EDEN-pes* is about 31% (*i.e.*, mobility traces that are always re-identified by FURIA). On the other extreme, *EDEN-opt* sends this portion of data without any protection or with a default LPPM. Roughly, 10% of mobility data is protected with the chosen LPPM while 20% of them are re-identified. However, we highlight that *EDEN-opt* ensures zero data loss. Finally, the balanced solution shows a tradeoff between data suppression and data publishing with a gain of +7% of protection and a raise of +3% of re-identification while the data loss is reduced to 21%.

Concerning the MDC dataset, the effect of EDEN's policies is not visible because the amount of data that FURIA is always able to re-identify is negligible (4%). In the Privamov dataset, the colors distribution is more balanced: around 20%, 11%, 8% and 3% of data are naturally protected, or protected with *PROM*, *TRL* and *Geoi*, respectively. However, the effect of EDEN's variants is the same as in the MDC dataset, with around 3% of re-identified mobility data.

A similar evaluation is done where the choice of NOBF is excluded and the default LPPM used to send data in case FURIA is always able to re-identify the mobility trace is *PROM-50m*. Results are depicted in Figure 5.12. Specifically, in Privamov and MDC datasets, LPPMs with low configurations (*TRL-1km, Geoi-0,01* and *PROM-50m*) replaces NOBF choice over EDEN's variants and the choice of *PROM-50m* as a default LPPM in EDEN-opt and EDEN-balanced does not impact the re-identification rate. However, in the Geolife dataset, the choice of *PROM-50m* to protect the 41% of deleted data in EDEN-pes increases the re-identification rate by +20% and the protection rate by +21% in EDEN-opt. In contrast, EDEN-balanced

registers only +5% of the former and +15% of the latter.

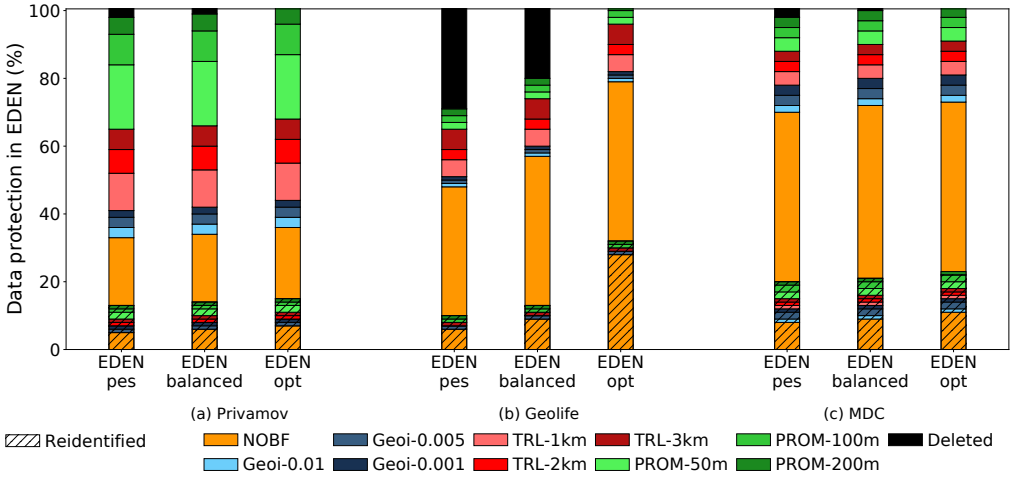


**Figure 5.11:** *Fine-grained analysis of* EDEN*'s variants.*

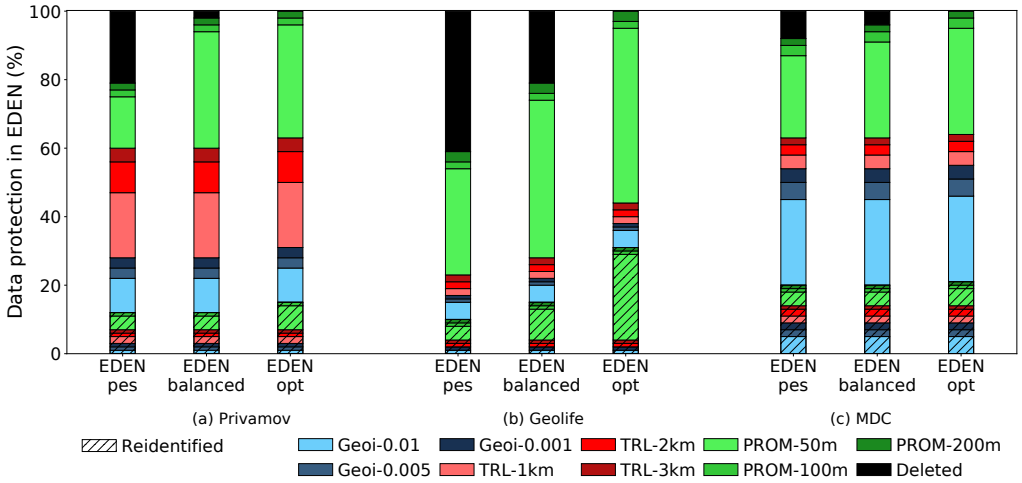


**Figure 5.12:** *Fine-grained analysis of* EDEN*'s variants without the NOBF choice.*

### 5.5.5 Evaluation of Performance Overhead

EDEN is a user-side protection approach that operates directly on edge devices. In this section, we measure the run-time induced by EDEN on different sizes of mobility data ranging from one mobility record, *i.e.*, equivalent to real-time crowd sensing applications, to longer mobility traces (up to 1,600 records). Table 5.3 provides more statistics about the considered datasets. The run-time overhead of EDEN

mainly includes the run-time for (1) the protection process where a set of LPPMs with different configurations are executed (*i.e.*, *Geoi, TRL, PROM*) and (2) the computation of AC metric to choose the best LPPM. In addition, the run-time for data processing (*i.e.*, converting data to feature vectors) and for the FURIA risk assessment are insignificant (nanoseconds). Results in Figure 5.13a with a logarithmic scale on the y-axis (for readability of small values) show that the longer a mobility trace, the longer it takes to protect it. Precisely, a mobility trace length ranging from 200 to 1,600 mobility records can take from 3 to 9 seconds to protect. This takes 15 milliseconds to protect a single mobility record which is acceptable in the context of real-time mobile crowd sensing applications. Moreover, we measure the execution time of the training phase in FURIA. Figure 5.13b illustrates the average training time per user in each learning round of different datasets. Overall datasets, we record a training time between 2 and 4 seconds per user. Specifically, in the Geolife dataset, the training time is slightly higher than Privamov and MDC, this is due to a higher number of participants in comparison to Privamov (*i.e.*, 19 users *vs.* 13 users) and denser mobility traces in comparison to MDC (*i.e.*, 323 records *vs.* 46 records in average). Our experiment is conducted on a desktop machine (see Section 5.5.1) and it is still practical while using a smartphone. For instance, the average time of applying an LPPM on an edge device is in order of milliseconds [44]. Also, the authors in [59] evaluate the computation time of a learning round on different smartphones. The latter is in order of seconds. Thus, EDEN which uses a simple, yet effective LR model can be deployed on real devices in future work.

**Table 5.3:** *Mobility dataset statistics.*

| Dataset | Average records per user | Standard deviation | Minimum records | Maximum records | Average #Users per round |
|---|---|---|---|---|---|
| Geolife | 323 | 288 | 1 | 1,800 | 19 |
| Privamov | 117 | 60 | 1 | 180 | 13 |
| MDC | 46 | 43 | 1 | 412 | 95 |

### 5.5.6 Quantifying the Fairness of EDEN

To push our evaluation further, we consider in this section another equally important aspect rarely addressed in the literature, that is, the concept of fairness [70]. It states

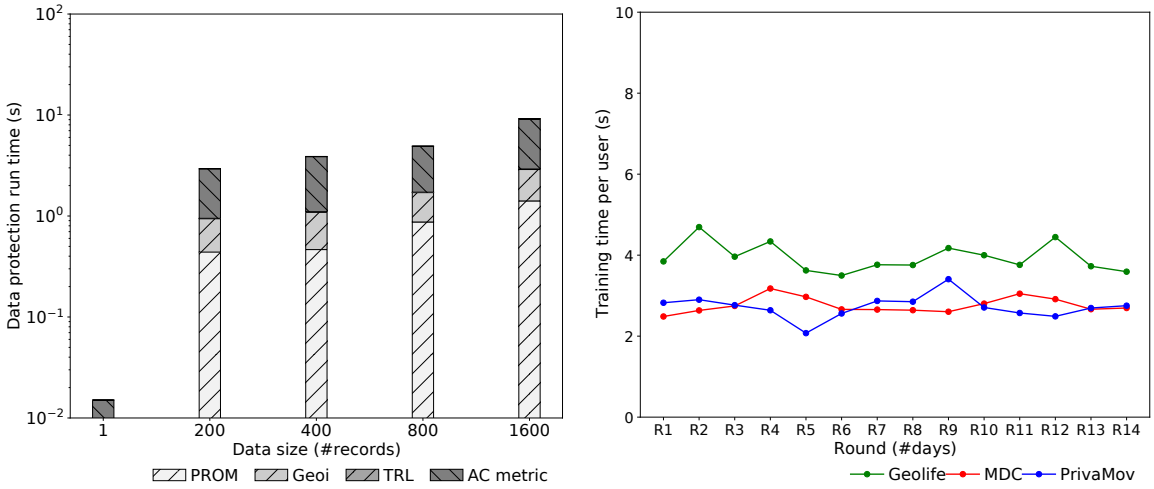


**(a)** EDEN *protection run-time* **(b)** *FURIA's model training time*

**Figure 5.13:** EDEN *run-time overhead*

that similar individuals should be treated similarly regarding their specific task [204]. In most of the cases, measuring the similarity between individuals is not a trivial task. In our context, we need two sets of definitions corresponding to the similarity between users' trajectories, and the similarity of the outcome of EDEN. We consider that individuals who are similar in terms of their mobility, should receive an equal privacy gain from EDEN. We define the similarity of trajectories by borrowing tenets from mobility literature and measure entropy of users as a measure of their maximum predictability. To this end, we use the Shannon Entropy (SE). The higher its value, the lower the predictability of an individual's movements [237, 146]. It is measured as in Equation 5.4.

$$E_h = -\sum_{i=1}^{n} P(x_i) \log_2[P(x_i)] \tag{5.4}$$

where $n$ is the length of probability vector, $P(x_i)$ is the probability of visiting location $x_i$ considering only the spatial pattern (*i.e.*, latitude and longitude).

In our evaluation, we hypothesize that users with similar entropy should receive similar privacy gain after applying EDEN. In an ideal setting we expect the entropy to increase for all the users (*i.e.*, predictability to decrease). Thus, we compare the entropy of similar user traces to their output by EDEN and we study in detail the percentage of users for whom the entropy decreases because of EDEN. We refer to

this group as the *disadvantaged group*

The results are depicted in Figure 5.14. We present the entropy before and after applying EDEN for Privamov, Geolife and MDC datasets. We observe that overall the datasets EDEN increases the entropy of most users. Indeed the cases where we find outcome of EDEN to *disadvantaged users* (decrease their entropy) are 8% for Privamov, 7% for Geolife, and 3% for MDC. Next, we study the fairness for those traces that correspond to the *disadvantaged group*.

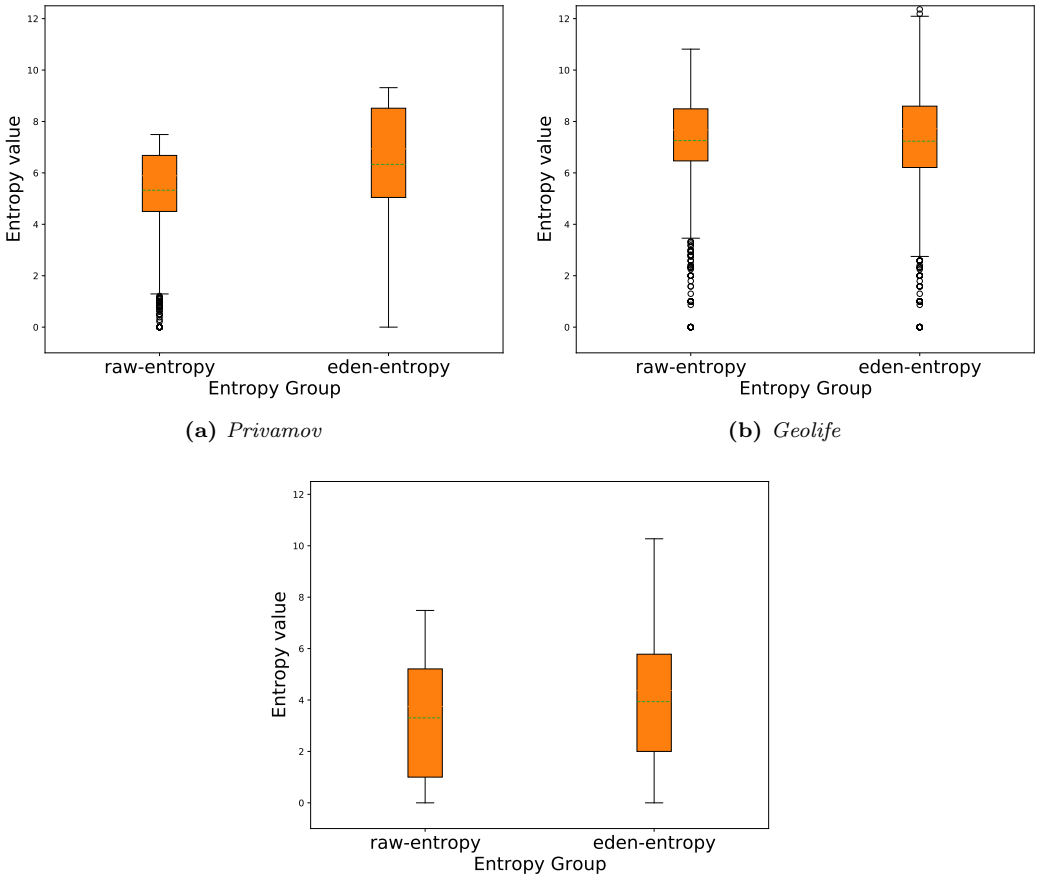


**(a)** *Privamov* **(b)** *Geolife*

**(c)** *MDC*

**Figure 5.14:** *Entropy level before and after the application of* EDEN.

Figure 5.15 illustrates the entropy decline for the *disadvantaged users* of all the datasets grouping them into entropy intervals. As we can see, the users with lower raw entropy (*i.e.*, prior to applying EDEN) receive a relatively less decline in their entropy after applying EDEN as well as smaller variations. In this plot, the size of

each box presents the fairness as measured by the difference in outcome after applying EDEN. Specifically, users who initially had lower predictability (high raw entropy) exhibit a larger variation in their entropy after applying EDEN, corresponding to different treatments (*i.e.*, unfair behavior of EDEN). Likewise, users with highly predictable patterns (low raw entropy) receive a similar outcome from EDEN (*i.e.*, fair behavior).

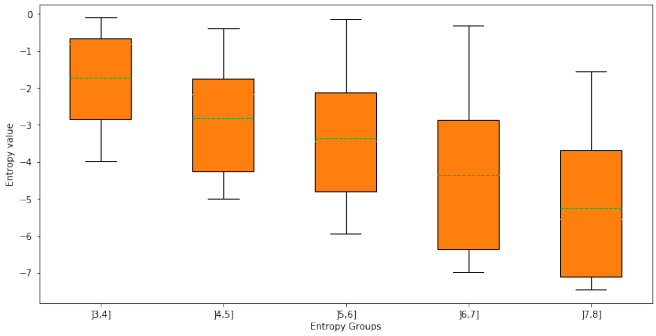


**(a)** *Privamov*

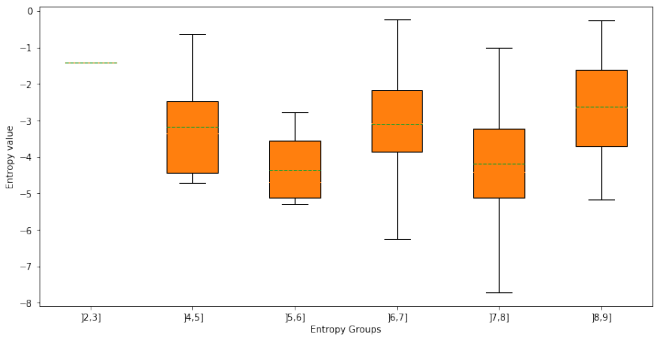


**(b)** *Geolife*

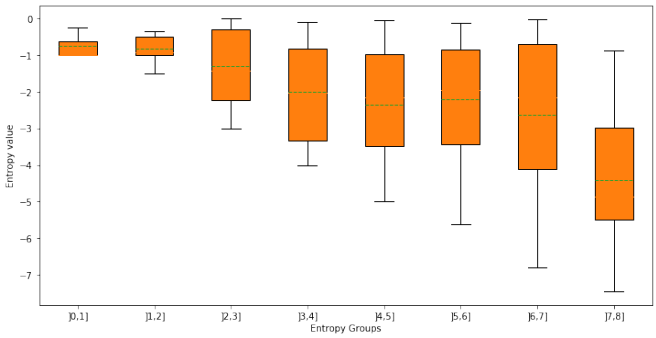


**(c)** *MDC*

**Figure 5.15:** *The entropy decline of disadvantaged groups.*

## 5.6 Privacy Discussion

**On the protection against compromised users.** EDEN does not protect against compromised users who might alter the FURIA model to reduce its performance. Nevertheless, the research community proposes several solutions to handle the issue of malicious users in the FL protocol with extensive countermeasures, presented in Chapter 2, Section 2.4.

**On the protection against a corrupted FURIA Master Server.** We assume that the FURIA Master Server is trusted. Nevertheless, there are solutions that protect against a compromised aggregator in the FL protocol. Secure aggregation protocols [39], differential privacy [93], and hardware solutions using trusted execution environments (TEEs) [164] are examples of these solutions and they are complementary to EDEN.

**On the protection against a malicious MCS.** In this work, EDEN aims to protect users' mobility data for crowd sensing applications where the MCS is honest-but-curious, *i.e.*, it faithfully provides the target service to its clients but tries to extract sensitive information from the received data. A harmful attacker may though take control of the MCS and perform arbitrary attacks such as publishing fake aggregate maps. This problem is orthogonal to EDEN and solutions such as using TEEs (*e.g.*, Intel SGX) can be considered to secure the computations performed on the MCS side.

**On the MCS running stronger attacks.** In this work, we considered *Mv-Attack* as the most powerful attack that the MCS may run to re-identify anonymous mobility data. However, the MCS may hold stronger attacks not published yet, leading to different results than the ones presented in this chapter. Nevertheless, as the MCS can become stronger by implementing new attacks, EDEN can also become stronger by implementing the latest LPPMs and their respective configurations.

**On the theoretical guarantees offered by EDEN.** EDEN builds on top of state-of-the-art LPPMs each of which comes with its own theoretical guarantees (*e.g.*, k-anonymity, differential privacy). EDEN adds on top of these guarantees a location privacy risk assessment, which allows choosing the most appropriate LPPM and its configuration to protect mobility traces according to the user re-identification risk and utility of the generated data.

## 5.7 Summary

In this chapter, we presented EDEN, a user-side protection system for mobility data in crowd sensing applications. It protects mobility data by choosing the best LPPM and configuration among a set of LPPMs without relying on a trusted proxy server. Instead, EDEN relies on the FL paradigm to train FURIA. It enables end-devices to train their local models using locally preserved mobility data while sharing the benefits of a global aggregated model across all users. This model is then used on the user's device to locally compare LPPMs for each mobility trace to protect and select the one which is resilient to FURIA. EDEN also relies on a utility metric to ensure a good quality of the resulting data.

We evaluated EDEN by performing a set of experiments on real-world mobility datasets. The results show that EDEN outperforms individual LPPMs both in terms of privacy measured by the resilience against a strong attack run by the MCS-side attacker and in terms of data utility measured using the AC metric. In addition, EDEN was also evaluated on a crowd sensed air pollution dataset. The results show that EDEN better preserves the distribution of gaseous pollutants compared to its competitors. Moreover, we also quantified the fairness of EDEN in our evaluation. The results show that EDEN has a fair behavior for almost the users (*i.e.*, increasing entropy of users' mobility and thus reducing their predictability). However, there is still a minority of users' mobility traces where EDEN fails to achieve fairness.

# Chapter 6

# Conclusion & Perspectives

## 6.1 Conclusion

In this thesis, we tackled the problem of location privacy in location-based services. First, we introduced some preliminaries and generalities about mobility data and how this resource is collected, stored, and processed by service providers. Specifically, an LBS is a double-edged weapon. On one side, it facilitates users' daily lives with a broad range of applications that provide personalized and customized information to users according to their location. But on the other side, a curious LBS or any entity that might have access to the gathered data may exploit it fraudulently to run several attacks and disclose sensitive and personal information about users. One of the most common threats highlighted in this thesis is the user re-identification risk. It aims to re-associate an anonymous mobility trace to its originating user based on historical mobility data. To overcome the privacy threats affecting mobility data, the research community has been actively proposing location privacy protection mechanisms (LPPMs). We defined LPPMs and classified them into three categories according to their use case scenarios. Specifically, there are online LPPMs for interactive and real-time use cases, semi-online LPPMs for crowd sensing campaigns, and offline LPPMs for data publishing use cases. These LPPMs transform mobility data with various techniques such as anonymization, perturbation, generalization, and fake data generation. To evaluate the effectiveness of these techniques, we presented three categories of metrics, namely privacy, utility and performance metrics. Moreover, we observed that each LPPM has its strengths and weaknesses, and the sensitivity of mobility data may differ from one user to another. That is why the protection of mobility data should be user-centric or even trace-centric. To reach this objective, authorities and organizations have strongly recommended conducting location privacy risk assessment (LPRA). The latter helps raising users' awareness and improving the protection by choosing the appropriate LPPM before sharing mobility data with service providers.

In this thesis, we proposed MOOD, **a centralized approach for location privacy protection** of data publishing use cases. It is a user-centric approach that aims to protect all users' mobility data, particularly orphan users who are not protected using individual LPPMs against re-identification attacks. It consists of two principles: a multi-LPPM composition where a set of LPPMs are combined sequentially and incrementally on a mobility trace and fine-grained protection where mobility traces are split into multiple sub-traces to separate discriminating mobility patterns. The obfuscation process is repeated until a stop criterion is verified. The latter can be the failure of the re-identification risk assessment (*i.e.*, the protected mobility trace or sub-trace is not re-identified) or the minimum length for mobility data. Finally, we demonstrated with extensive experiments that MOOD effectively protects mobility datasets with a good privacy *vs.* utility tradeoff.

Furthermore, existing works including MOOD rely on a centralized trusted server to carry out both the obfuscation process and the location privacy risk assessment. This constitutes a single point of failure because, if the server is compromised, mobility data can be leaked and users' privacy may be threatened. To avoid centralizing data, we proposed SAFER, **a federated approach to assess the privacy risk of sharing mobility data** while keeping it on the user's premises. SAFER evolves a federated identity classifier and a uniqueness evaluation. The federated identity classifier is trained on mobility data using the FL principles to build a global knowledge about users' mobility behavior. It produces a confidence vector that contains the probabilities of owning the considered mobility data by different users in the system. Then, to evaluate the uniqueness of that data, the confidence vector allows constructing anonymity sets by grouping users whose probabilities are close. The size of the anonymity set reflects how unique mobility data is and thus can quantify the privacy bounds of human mobility in a privacy-preserving way with scalable and comparable results to that of a well-established baseline system.

Finally, **to enforce location privacy without centralizing raw mobility data, we proposed EDEN**. In the latter, the protection process and the location privacy risk assessment are performed on the edge device thanks to the FL paradigm. Specifically, EDEN goes through two phases. The first phase consists of a re-identification risk assessment where a federated model called FURIA is periodically trained on the users' devices to learn the mobility patterns of each user (*i.e.*, class). The second phase is a protection phase where the latest aggregate model is used

along with utility metrics to automatically choose the appropriate LPPM with its corresponding configuration each time a user wants to send a geo-located trace to the service provider. We evaluated EDEN on five real-word mobility datasets and demonstrated its effectiveness in terms of privacy, utility and performance.

## 6.2 Perspectives

By the end of this thesis, many areas and perspectives still need to be explored in the literature. We consider both short and long-term perspectives.

### 6.2.1 Short-term Perspectives

Although we believe that our contributions can improve location privacy with effective results, we are aware that there is still room for improvement. We discuss in the following the main possible improvements that we plan to consider shortly.

#### Improving Data Protection with MOOD

MOOD is a promising solution for the research community on location privacy. It can be extended with the latest state-of-the-art LPPMs, attacks, and utility metrics. However, increasing the number of LPPMs and attacks is time-consuming as MOOD is based on a brute force search to find the most optimal composition of LPPMs (the one that passes the re-identification risk assessment and maintains the highest data utility). Fortunately, MOOD concerns data publishing use cases where time is not a constraint. However, the application of MOOD in other use cases such as interactive or crowd sensing requires optimizing the search by exploring machine learning techniques to build predictive models. The latter would automatically select the most appropriate LPPM in order to avoid the exhaustive search. In addition, MOOD follows fine-grained data protection. It splits mobility traces into sub-traces using fixed time slices. In this direction, our objective is to explore other relevant ways of data splitting, such as considering the semantics of visited locations, the gap between them *etc.*,.

### Assessing Uniqueness with SAFER

SAFER is a user-side, location privacy risk assessment. It is suitable for interactive use cases where an end user evaluates the uniqueness of a given data point in real-time before leaving the mobile device. It exploits the latest version of the federated model, generated at night to evaluate the uniqueness of mobility data. In the future, we should investigate how to adapt the training frequency of the global model according to the performance of users' devices in a realistic FL environment. Specifically, mobile devices with low performance should continue to train at night, while those with high performance could update the model whenever they wish to assess privacy risks. In addition, SAFER is currently applied to raw mobility data. As a future work, we should study the impact of SAFER when mobility data is obfuscated with protection mechanisms, such as perturbation, generalization, and fake data generation. Furthermore, the current work, as previously mentioned in Section 4.5, does not consider compromised users and a malicious aggregator in the FL protocol. In this direction, SAFER should be implemented using TEEs in an end-to-end manner both on the client side and on the server side. In this way, all the exchanged messages are encrypted and they are only used in clear inside enclaves.

### Protecting Mobility Data with EDEN

EDEN is based on a simple yet effective location privacy risk assessment model called FURIA. In the future, we should investigate advanced models to train FURIA without the feature engineering step which might lack relevant information that discriminates human mobility. In this direction, recurrent networks is a good starting point, previously used with effective results for the uniqueness assessment. The latter captures recurrent moving patterns in human mobility, especially if we manipulate long trajectories. In addition, the architecture of FURIA requires a trusted master server that might infer sensitive information about users. That is why it is necessary to implement and combine security primitives, such as secure aggregation, differential privacy, or even secure hardware (*e.g.*, TEEs) to prevent any information leakage or inference attacks. Specifically, EDEN could be combined with other state-of-the-art orthogonal techniques in order to face compromised users [23, 223, 34], or to counter model information leakage attacks [39, 93, 164].

### 6.2.2 Long-term Perspectives

The research area on location privacy has known unprecedented growth in the last decade. The progress of LPPMs is in perpetual competition with the development of novel attacks that try to break these mechanisms along with the advance in information and communication technologies. That is why there is still a long way to go until location privacy can be democratized and controlled politically by raising people's awareness and technically by producing software where privacy requirements are integrated by design and default. Specifically, people become highly dependent on geo-located services. They can freely share their mobility data at the cost of their privacy. This is due to users not being sufficiently aware of the value of their mobility data and the amount of sensitive knowledge that can be derived from it. In this sense, it is essential to increase users' awareness of the sensitivity of their mobility by proposing intuitive and convenient tools to highlight privacy issues and the benefits of using LPPMs. For instance, Please Rob Me[1] is a website to raise users' awareness about what can be revealed from shared geolocated tagged tweets. The latter may inform potential robbers about the presence or absence of a target user at home.

On the other side of the spectrum, companies that manipulate mobility data do not miss out on the opportunity to infer sensitive information about users from the collected data or sell it to third parties for unauthorized purposes. Hence, legislation and regulations are likely needed to prevent users' privacy from being violated. This is a fundamental right that companies should respect with complete transparency about the usage of the gathered data. In this thesis, we encourage and enforce these laws by proposing practical systems to better preserve users' privacy while maintaining a good data utility. These solutions might be used by companies and practitioners to enhance their privacy policies in compliance with the GDPR's recommendations.

In such a way, existing LPPMs are already used in real-life scenarios such as Geo-indistinguishability [21] which has been implemented by its authors as a browser extension [9]. This allows users to benefit from some privacy when using LBSs through their web browser. In this context, it would be interesting to deploy SAFER or EDEN in a real-life environment using libraries and frameworks to implement LPPMs and machine learning models on edge devices. For that purpose, Android

[1] https://pleaserobme.com/

supports a wide variety of helpful tools such as python, ML-Kit[2], *etc.*,.

However, there are still challenges to make these solutions suitable for real-life use cases. For instance, SAFER uses a fixed number of classes in the federated identity classifier. In the future, it would be interesting to investigate how to build incremental online models, which are able to handle larger numbers of incoming users (order of millions) without forgetting their past knowledge, previously learned. In addition, with the growth of collected mobility data generated by users' devices, a completely decentralized learning model (*i.e.*, gossip learning model [113]) may be a promising alternative to federated learning as the latter generally relies on a central entity that may engender bottlenecks, inhibit scalability of the system, and cause privacy issues.

Last and not least, a challenging research direction is to develop learning methods that combine different data modalities. Specifically, the majority of existing studies in the literature assume a single data modality, *e.g.*, trajectories, geolocated social media data (posts, pictures, or videos). In reality, data come in several modalities from different sources, and combining those can be used to extract richer representations and build more accurate models. However, it may open the door to more privacy threats in the future. For instance, a user requesting a target destination on Google maps application using her microphone (voice data) combined with her location may reveal more sensitive information about the user (*e.g.*, gender, physical condition, emotions, *etc.*,).

[2]https://developers.google.com/ml-kit

# List of Figures

# List of Tables

# Bibliography


[1] MCOC-Shanghai – A Dataset From a Major Cellular Operator in China. Non-public dataset. 17, 69, 79

[2] Core ML. URL https://developer.apple.com/machine-learning/core-ml/. 93

[3] 533 Million Facebook Users' Phone Numbers and Personal Data Have Been Leaked Online, . URL https://www.businessinsider.com/stolen-data-of-533-million-facebook-users-leaked-online-2021-4. 68

[4] Facebook. Facebook Places, . URL https://developers.facebook.com/docs/places/. 18

[5] Manage screen & display settings on a Pixel phone. URL https://support.google.com/pixelphone/answer/6111557. 93

[6] Defending the NHS from Cyber-Attacks in 2022. URL https://healthcareglobal.com/hospitals/defending-the-nhs-from-cyber-attacks-in-2022. 68

[7] Keras library. URL https://www.keras.io. 77

[8] LinkedIn Data Breach 2021. URL https://economictimes.indiatimes.com/tech/technology/linkedin-denies-data-breach-that-reportedly-affected-700-million-users/articleshow/83977790.cms. 68

[9] Location guard. URL https://github.com/chatziko/location-guard. 137

[10] Pytorch library. URL https://www.pytorch.org. 77, 111

[11] S2 geometry. URL https://www.s2geometry.io. 77, 111

[12] Google maps. https://www.google.fr/maps, 2005. 4

[13] Gboard - the google keyboard, 2016. URL https://apps.apple.com/fr/app/gboard-le-clavier-google/id1091700242. 93

[14] Osman Abul, Francesco Bonchi, and Mirco Nanni. Never walk alone: Uncertainty for anonymity in moving objects databases. In *2008 IEEE 24th international conference on data engineering*, pages 376–385. Ieee, 2008. 27

[15] Osman Abul, Francesco Bonchi, and Mirco Nanni. Anonymization of moving objects databases by clustering and perturbation. *Information systems*, 35(8): 884–910, 2010. 27, 114

[16] Naman Agarwal, Ananda Theertha Suresh, Felix Xinnan X Yu, Sanjiv Kumar, and Brendan McMahan. cpsgd: Communication-efficient and differentially-private distributed sgd. *Advances in Neural Information Processing Systems*, 31, 2018. 95

[17] Ulrich Matchi Aïvodji, Kévin Huguenin, Marie-José Huguet, and Marc-Olivier Killijian. Sride: A privacy-preserving ridesharing system. In *Proceedings of the 11th ACM Conference on Security & Privacy in Wireless and Mobile Networks*, pages 40–50, 2018. 31

[18] Zaheer Allam and David S Jones. On the coronavirus (covid-19) outbreak and the smart city network: universal data sharing standards coupled with artificial intelligence (ai) to benefit urban health monitoring and management. In *Healthcare*, volume 8, page 46. MDPI, 2020. 4

[19] Saed Alrabaee, Noman Saleem, Stere Preda, Lingyu Wang, and Mourad Debbabi. Corrigendum to 'oba2: An onion approach to binary code authorship attribution' [digit investig 11 (2014) S94-S103]. *Digit. Investig.*, 21:89, 2017. doi: 10.1016/j.diin.2017.02.004. URL https://doi.org/10.1016/j.diin.2017.02.004. 21

[20] Sebastien Andreina, Giorgia Azzurra Marson, Helen Möllering, and Ghassan Karame. Baffle: Backdoor detection via feedback-based federated learning. In *2021 IEEE 41st International Conference on Distributed Computing Systems (ICDCS)*, pages 852–863. IEEE, 2021. 42, 94

[21] Miguel E Andrés, Nicolás E Bordenabe, Konstantinos Chatzikokolakis, and Catuscia Palamidessi. Geo-indistinguishability: Differential privacy for location-based systems. In *Proceedings of the 2013 ACM SIGSAC conference on Computer & communications security*, pages 901–914, 2013. 98, 101, 115, 137

[22] Miguel E. Andrés, Nicolás Emilio Bordenabe, Konstantinos Chatzikokolakis, and Catuscia Palamidessi. Geo-indistinguishability: differential privacy for location-based systems. In *2013 ACM SIGSAC Conference on Computer and Communications Security, CCS'13, Berlin, Germany, November 4-8, 2013*, pages 901–914, 2013. doi: 10.1145/2508859.2516735. URL https://doi.org/10.1145/2508859.2516735. 29, 31, 33, 45, 56

[23] Eugene Bagdasaryan, Andreas Veit, Yiqing Hua, Deborah Estrin, and Vitaly Shmatikov. How to Backdoor Federated Learning. In *International Conference on Artificial Intelligence and Statistics*, pages 2938–2948. PMLR, 2020. 40, 94, 136

[24] Daniel Balouek, Alexandra Carpen Amarie, Ghislain Charrier, Frédéric Desprez, Emmanuel Jeannot, Emmanuel Jeanvoine, Adrien Lèbre, David Margery, Nicolas Niclausse, Lucas Nussbaum, Olivier Richard, Christian Pérez, Flavien Quesnel, Cyril Rohr, and Luc Sarzyniec. Adding Virtualization Capabilities to the Grid'5000 Testbed. In Ivan I. Ivanov, Marten van Sinderen, Frank Leymann, and Tony Shan, editors, *Cloud Computing and Services Science*, volume 367 of *Communications in Computer and Information Science*, pages 3–20. Springer International Publishing, 2013. ISBN 978-3-319-04518-4. doi: 10.1007/978-3-319-04519-1_1. 78

[25] Omer Barak, Gabriella Cohen, and Eran Toch. Anonymizing mobility data using semantic cloaking. *Pervasive and Mobile Computing*, 28:102–112, 2016. 28

[26] Marco Barreno, Blaine Nelson, Russell Sears, Anthony D Joseph, and J Doug Tygar. Can machine learning be secure? In *Proceedings of the 2006 ACM Symposium on Information, computer and communications security*, pages 16–25, 2006. 94

[27] Anahid Basiri, Terry Moore, and Chris Hill. The privacy concerns in location based services: protection approaches and remaining challenges. 2016. 6

[28] Christine Bauer and Christine Strauss. Location-based advertising on mobile devices. *Management review quarterly*, 66(3):159–194, 2016. 4

[29] Mihir Bellare, Oded Goldreich, and Shafi Goldwasser. Incremental cryptography: The case of hashing and signing. In *Annual International Cryptology Conference*, pages 216–233. Springer, 1994. 95

[30] Alastair R Beresford and Frank Stajano. Mix zones: User privacy in location-aware services. In *IEEE Annual conference on pervasive computing and communications workshops, 2004. Proceedings of the Second*, pages 127–131. IEEE, 2004. 26

[31] Christophe Bertero. *Perception of the urban environment with the help of a fleet of sensors on bicycles. Application to air pollution.* PhD thesis, Toulouse University, Toulouse University, France, 2020. 17, 99, 113, 117, 120

[32] Claudio Bettini, X Sean Wang, and Sushil Jajodia. Protecting privacy against location-based personal identification. In *Workshop on Secure Data Management*, pages 185–199. Springer, 2005. 36

[33] Claudio Bettini, Gabriele Civitarese, and Riccardo Presotto. Personalized semi-supervised federated learning for human activity recognition. *arXiv preprint arXiv:2104.08094*, 2021. 71

[34] Arjun Nitin Bhagoji, Supriyo Chakraborty, Prateek Mittal, and Seraphin Calo. Analyzing federated learning through an adversarial lens. In *International Conference on Machine Learning*, pages 634–643. PMLR, 2019. 40, 94, 136

[35] Igor Bilogrevic, Kévin Huguenin, Murtuza Jadliwala, Florent Lopez, Jean-Pierre Hubaux, Philip Ginzboorg, and Valtteri Niemi. Inferring social ties in academic networks using short-range wireless communications. In *Proceedings of the 12th ACM workshop on Workshop on privacy in the electronic society*, pages 179–188. ACM, 2013. 5, 19

[36] Vincent Bindschaedler and Reza Shokri. Synthesizing plausible privacy-preserving location traces. In *2016 IEEE Symposium on Security and Privacy (SP)*, pages 546–563. IEEE, 2016. 28

[37] Peva Blanchard, El Mahdi El Mhamdi, Rachid Guerraoui, and Julien Stainer. Machine learning with adversaries: Byzantine tolerant gradient descent. *Advances in Neural Information Processing Systems*, 30, 2017. 42, 94

[38] Hristo Bojinov, Yan Michalevsky, Gabi Nakibly, and Dan Boneh. Mobile Device Identification via Sensor Fingerprinting. *CoRR*, abs/1408.1416, 2014. URL http://arxiv.org/abs/1408.1416. 36

[39] Keith Bonawitz, Vladimir Ivanov, Ben Kreuter, Antonio Marcedone, H Brendan McMahan, Sarvar Patel, Daniel Ramage, Aaron Segal, and Karn Seth. Practical secure aggregation for privacy-preserving machine learning. In *Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security*, pages 1175–1191, 2017. 41, 71, 72, 131, 136

[40] Keith Bonawitz, Hubert Eichner, Wolfgang Grieskamp, Dzmitry Huba, Alex Ingerman, Vladimir Ivanov, Chloe Kiddon, Jakub Konečnỳ, Stefano Mazzocchi, Brendan McMahan, et al. Towards federated learning at scale: System design. *Proceedings of Machine Learning and Systems*, 1:374–388, 2019. 71

[41] Guillaume Bouchard. Efficient bounds for the softmax function, applications to inference in hybrid models. In *Presentation at the Workshop for Approximate Bayesian Inference in Continuous/Hybrid Systems at NIPS-07*. Citeseer, 2007. 75

[42] Spyros Boukoros, Mathias Humbert, Stefan Katzenbeisser, and Carmela Troncoso. On (the lack of) location privacy in crowdsourcing applications. In Nadia Heninger and Patrick Traynor, editors, *28th USENIX Security Symposium, USENIX Security 2019, Santa Clara, CA, USA, August 14-16, 2019*, pages 1859–1876. USENIX Association, 2019. URL https://www.usenix.org/conference/usenixsecurity19/presentation/boukoros. 34

[43] Antoine Boutet and Sonia Ben Mokhtar. Uniqueness Assessment of Human Mobility on Multi-Sensor Datasets. In *The 16th International Conference on Availability, Reliability and Security*, pages 1–10, 2021. 36, 68

[44] Ioannis Boutsis and Vana Kalogeraki. Location privacy for crowdsourcing applications. In *Proceedings of the 2016 ACM International Joint Conference on Pervasive and Ubiquitous Computing*, pages 694–705, 2016. 127

[45] Leo Breiman. Random forests. *Machine learning*, 45(1):5–32, 2001. 78, 105

[46] Azby Brown, Pieter Franken, Sean Bonner, Nick Dolezal, and Joe Moross. Safecast: successful citizen-science for radiation measurement and communication after fukushima. *Journal of Radiological Protection*, 36(2):S82, 2016. 4, 17, 98

[47] Aylin Caliskan, Fabian Yamaguchi, Edwin Dauber, Richard E. Harang, Konrad Rieck, Rachel Greenstadt, and Arvind Narayanan. When coding style survives compilation: De-anonymizing programmers from executable binaries. In *25th Annual Network and Distributed System Security Symposium, NDSS 2018, San Diego, California, USA, February 18-21, 2018*. The Internet Society, 2018. URL http://wp.internetsociety.org/ndss/wp-content/uploads/sites/25/2018/02/ndss2018_06B-2_Caliskan_paper.pdf. 21

[48] Hancheng Cao, Jie Feng, Yong Li, and Vassilis Kostakos. Uniqueness in the City: Urban Morphology and Location Privacy. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies*, 2(2):1–20, 2018. 36

[49] Alket Cecaj, Marco Mamei, and Franco Zambonelli. Re-identification and information fusion between anonymized cdr and social network data. *Journal of Ambient Intelligence and Humanized Computing*, 7(1):83–96, 2016. 23

[50] Sophie Cerf, Vincent Primault, Antoine Boutet, Sonia Ben Mokhtar, Robert Birke, Sara Bouchenak, Lydia Y. Chen, Nicolas Marchand, and Bogdan Robu. PULP: achieving privacy and utility trade-off in user mobility data. In *36th IEEE Symposium on Reliable Distributed Systems, SRDS 2017, Hong Kong, Hong Kong, September 26-29, 2017*, pages 164–173. IEEE Computer Society, 2017. doi: 10.1109/SRDS.2017.25. URL https://doi.org/10.1109/SRDS.2017.25. 56

[51] Sophie Cerf, Vincent Primault, Antoine Boutet, Sonia Ben Mokhtar, Robert Birke, Sara Bouchenak, Lydia Y Chen, Nicolas Marchand, and Bogdan Robu. Pulp: achieving privacy and utility trade-off in user mobility data. In *2017 IEEE 36th Symposium on Reliable Distributed Systems (SRDS)*, pages 164–173. IEEE, 2017. 32

[52] Sung-Hyuk Cha. Comprehensive survey on distance/similarity measures between probability density functions. *City*, 1(2):1, 2007. 109

[53] Konstantinos Chatzikokolakis, Catuscia Palamidessi, and Marco Stronati. Constructing elastic distinguishability metrics for location privacy. *Proc. Priv. Enhancing Technol.*, 2015(2):156–170, 2015. doi: 10.1515/popets-2015-0023. URL https://doi.org/10.1515/popets-2015-0023. 29

[54] Meng Chen, Yang Liu, and Xiaohui Yu. Nlpmm: A next location predictor with markov modeling. In *Pacific-Asia conference on knowledge discovery and data mining*, pages 186–197. Springer, 2014. 5, 20

[55] Mingqing Chen, Rajiv Mathews, Tom Ouyang, and Françoise Beaufays. Federated learning of out-of-vocabulary words. *arXiv preprint arXiv:1903.10635*, 2019. 38

[56] Zhenyu Chen, Yanyan Fu, Min Zhang, Zhenfeng Zhang, and Hao Li. The de-anonymization method based on user spatio-temporal mobility trace. In *International Conference on Information and Communications Security*, pages 459–471. Springer, 2017. 22

[57] Benny Chor, Oded Goldreich, Eyal Kushilevitz, and Madhu Sudan. Private information retrieval. In *Proceedings of IEEE 36th Annual Foundations of Computer Science*, pages 41–50. IEEE, 1995. 30

[58] Chi-Yin Chow, Mohamed F Mokbel, and Xuan Liu. A peer-to-peer spatial cloaking algorithm for anonymous location-based service. In *Proceedings of the 14th annual ACM international symposium on Advances in geographic information systems*, pages 171–178, 2006. 27, 37

[59] Georgios Damaskinos, Rachid Guerraoui, Anne-Marie Kermarrec, Vlad Nitu, Rhicheek Patra, and François Taïani. Fleet: Online federated learning via staleness awareness and performance prediction. In Dilma Da Silva and Rüdiger Kapitza, editors, *Middleware '20: 21st International Middleware Conference, Delft, The Netherlands, December 7-11, 2020*, pages 163–177. ACM, 2020. doi: 10.1145/3423211.3425685. URL https://doi.org/10.1145/3423211.3425685. 127

[60] Georgios Damaskinos, Rachid Guerraoui, Anne-Marie Kermarrec, Vlad Nitu, Rhicheek Patra, and Francois Taiani. Fleet: Online federated learning via staleness awareness and performance prediction. In *Proceedings of the 21st International Middleware Conference*, pages 163–177, 2020. 71

[61] Weizhen Dang, Haibo Wang, Shirui Pan, Pei Zhang, Chuan Zhou, Xin Chen, and Jilong Wang. Predicting human mobility via graph convolutional dual-attentive networks. In *Proceedings of the Fifteenth ACM International Conference on Web Search and Data Mining*, pages 192–200, 2022. 20

[62] Fida Kamal Dankar, Khaled El Emam, Angelica Neisa, and Tyson Roffey. Estimating the re-identification risk of clinical data sets. *BMC Medical Informatics Decis. Mak.*, 12:66, 2012. doi: 10.1186/1472-6947-12-66. URL https://doi.org/10.1186/1472-6947-12-66. 21

[63] Commission Nationale de l'informatique et des Liberte (CNIL). Privacy impact assessment (pia) methodology—how to carry out a pia. 2015. 35

[64] Yves-Alexandre De Montjoye, César A Hidalgo, Michel Verleysen, and Vincent D Blondel. Unique in the Crowd: The Privacy Bounds of Human Mobility. *Scientific Reports*, 3(1):1–5, 2013. 5, 12, 26, 36, 68, 69, 79, 85

[65] Mina Deng, Kim Wuyts, Riccardo Scandariato, Bart Preneel, and Wouter Joosen. A privacy threat analysis framework: supporting the elicitation and fulfillment of privacy requirements. *Requirements Engineering*, 16(1):3–32, 2011. 35

[66] Srinivas Devarakonda, Parveen Sevusu, Hongzhang Liu, Ruilin Liu, Liviu Iftode, and Badri Nath. Real-time air quality monitoring through mobile sensing in metropolitan areas. In *Proceedings of the 2nd ACM SIGKDD international workshop on urban computing*, pages 1–8, 2013. 98

[67] Kai Dong, Taolin Guo, Haibo Ye, Xuansong Li, and Zhen Ling. On the limitations of existing notions of location privacy. *Future Generation Computer Systems*, 86:1513–1522, 2018. 33

[68] Prabal Dutta, Paul M Aoki, Neil Kumar, Alan Mainwaring, Chris Myers, Wesley Willett, and Allison Woodruff. Common sense: participatory urban sensing using a network of handheld air quality monitors. In *Proceedings of the 7th ACM conference on embedded networked sensor systems*, pages 349–350, 2009. 17

[69] Cynthia Dwork. Differential privacy: A survey of results. In *International conference on theory and applications of models of computation*, pages 1–19. Springer, 2008. 29, 33, 41, 98

[70] Cynthia Dwork, Moritz Hardt, Toniann Pitassi, Omer Reingold, and Richard Zemel. Fairness through awareness. In *Proceedings of the 3rd innovations in theoretical computer science conference*, pages 214–226, 2012. 127

[71] Peter Eckersley. How Unique Is Your Web Browser? In Mikhail J. Atallah and Nicholas J. Hopper, editors, *Privacy Enhancing Technologies, 10th International Symposium, PETS 2010, Berlin, Germany, July 21-23, 2010. Proceedings*, volume 6205 of *Lecture Notes in Computer Science*, pages 1–18. Springer, 2010. doi: 10.1007/978-3-642-14527-8_1. URL https://doi.org/10.1007/978-3-642-14527-8_1. 36

[72] Dominik Maria Endres and Johannes E Schindelin. A new metric for probability distributions. *IEEE Transactions on Information theory*, 2003. 22

[73] Minghong Fang, Xiaoyu Cao, Jinyuan Jia, and Neil Gong. Local model poisoning attacks to {Byzantine-Robust} federated learning. In *29th USENIX Security Symposium (USENIX Security 20)*, pages 1605–1622, 2020. 40

[74] Kassem Fawaz and Kang G Shin. Location privacy protection for smartphone users. In *Proceedings of the 2014 ACM SIGSAC Conference on Computer and Communications Security*, pages 239–250, 2014. 32

[75] Jie Feng, Mingyang Zhang, Huandong Wang, Zeyu Yang, Chao Zhang, Yong Li, and Depeng Jin. Dplink: User identity linkage via deep neural network from heterogeneous mobility data. In *The World Wide Web Conference*, pages 459–469, 2019. 26

[76] Jie Feng, Yong Li, Zeyu Yang, Qiang Qiu, and Depeng Jin. Predicting human mobility with semantic motivation via multi-task attentional recurrent networks. *IEEE Transactions on Knowledge and Data Engineering*, 2020. 20

[77] Jie Feng, Can Rong, Funing Sun, Diansheng Guo, and Yong Li. Pmf: A privacy-preserving human mobility prediction framework via federated learning. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies*, 4(1):1–21, 2020. 20, 41

[78] David Förster. Decentralized enforcement of k-anonymity for location privacy using secret sharing. In *Verifiable Privacy Protection for Vehicular Communication Systems*, pages 93–121. Springer, 2017. 37

[79] Lorenzo Franceschi-Bicchierai. Edditor Cracks Anonymous Data Trove to Pinpoint Muslim Cab Drivers. *Online at: http://mashable.com/2015/01/28/redditor-muslim-cab-drivers*, 2015. 5, 18

[80] Matteo Francia, Enrico Gallinucci, Matteo Golfarelli, and Nicola Santolini. Dart: De-anonymization of personal gazetteers through social trajectories. *Journal of Information Security and Applications*, 55:102634, 2020. 23

[81] Julien Freudiger, Reza Shokri, and Jean-Pierre Hubaux. On the optimal placement of mix zones. In *International Symposium on Privacy Enhancing Technologies Symposium*, pages 216–234. Springer, 2009. 26

[82] Nicholas Frosst and Geoffrey Hinton. Distilling a Neural Network into a Soft Decision Tree. *arXiv preprint arXiv:1711.09784*. 78

[83] Clement Fung, Chris JM Yoon, and Ivan Beschastnikh. Mitigating Sybils in Federated Learning Poisoning. *arXiv preprint arXiv:1808.04866*, 2018. 42, 94

[84] Sébastien Gambs, Marc-Olivier Killijian, and Miguel Núñez del Prado Cortez. GEPETO: A geoprivacy-enhancing toolkit. In *24th IEEE International Conference on Advanced Information Networking and Applications Workshops, WAINA 2010, Perth, Australia, 20-13 April 2010*, pages 1071–1076. IEEE Computer Society, 2010. doi: 10.1109/WAINA.2010.170. URL https://doi.org/10.1109/WAINA.2010.170. 18

[85] Sébastien Gambs, Marc-Olivier Killijian, and Miguel Núñez del Prado Cortez. Show me how you move and i will tell you who you are. In *Proceedings of the 3rd ACM SIGSPATIAL International Workshop on Security and Privacy in GIS and LBS*, pages 34–41, 2010. 18

[86] Sébastien Gambs, Marc-Olivier Killijian, and Miguel Núñez del Prado Cortez. Next place prediction using mobility markov chains. In *Proceedings of the First Workshop on Measurement, Privacy, and Mobility*, page 3. ACM, 2012. 20

[87] Sébastien Gambs, Marc-Olivier Killijian, and Miguel Núñez del Prado Cortez. De-anonymization attack on geolocated data. In *12th IEEE International Conference on Trust, Security and Privacy in Computing and Communications, TrustCom 2013 / 11th IEEE International Symposium on Parallel and Distributed Processing with Applications, ISPA-13 / 12th IEEE International Conference on Ubiquitous Computing and Communications, IUCC-2013, Melbourne, Australia, July 16-18, 2013*, pages 789–797. IEEE Computer Society, 2013. doi: 10.1109/TrustCom.2013.96. URL https://doi.org/10.1109/TrustCom.2013.96. 22, 45, 98, 99, 105, 115

[88] Karan Ganju, Qi Wang, Wei Yang, Carl A Gunter, and Nikita Borisov. Property inference attacks on fully connected neural networks using permutation invariant representations. In *Proceedings of the 2018 ACM SIGSAC conference on computer and communications security*, pages 619–633, 2018. 41

[89] Raghu K Ganti, Fan Ye, and Hui Lei. Mobile crowdsensing: current state and future challenges. *IEEE communications Magazine*, 49(11):32–39, 2011. 17, 112

[90] Qiang Gao, Fan Zhou, Kunpeng Zhang, Goce Trajcevski, Xucheng Luo, and Fengli Zhang. Identifying Human Mobility via Trajectory Embeddings. In *Proceedings of the 26th International Joint Conference on Artificial Intelligence (IJCAI 2017)*, volume 17, pages 1689–1695, 2017. 74, 75

[91] Bugra Gedik and Ling Liu. Location privacy in mobile systems: A personalized anonymization model. In *25th IEEE International Conference on Distributed Computing Systems (ICDCS'05)*, pages 620–629. IEEE, 2005. 27

[92] Arthur Gervais, Reza Shokri, Adish Singla, Srdjan Capkun, and Vincent Lenders. Quantifying web-search privacy. In Gail-Joon Ahn, Moti Yung, and Ninghui Li, editors, *Proceedings of the 2014 ACM SIGSAC Conference on Computer and Communications Security, Scottsdale, AZ, USA, November 3-7, 2014*, pages 966–977. ACM, 2014. doi: 10.1145/2660267.2660367. URL https://doi.org/10.1145/2660267.2660367. 21

[93] Robin C Geyer, Tassilo Klein, and Moin Nabi. Differentially private federated learning: A client level perspective. *arXiv preprint arXiv:1712.07557*, 2017. 131, 136

[94] Meysam Ghaffari, Nasser Ghadiri, Mohammad Hossein Manshaei, and Mehran Sadeghi Lahijani. A peer-to-peer privacy preserving query service for location-based mobile applications. *IEEE Transactions on Vehicular Technology*, 66(10):9458–9469, 2017. 37

[95] Gabriel Ghinita, Panos Kalnis, and Spiros Skiadopoulos. Mobihide: a mobile a peer-to-peer system for anonymous location-based queries. In *International Symposium on Spatial and Temporal Databases*, pages 221–238. Springer, 2007. 37

[96] Gabriel Ghinita, Panos Kalnis, and Spiros Skiadopoulos. Prive: anonymous location-based queries in distributed mobile systems. In *Proceedings of the 16th international conference on World Wide Web*, pages 371–380, 2007. 37

[97] Gabriel Ghinita, Panos Kalnis, Ali Khoshgozaran, Cyrus Shahabi, and Kian-Lee Tan. Private queries in location based services: anonymizers are not necessary. In *Proceedings of the 2008 ACM SIGMOD international conference on Management of data*, pages 121–132, 2008. 30

[98] Gabriel Ghinita, Maria Luisa Damiani, Claudio Silvestri, and Elisa Bertino. Preventing velocity-based linkage attacks in location-aware applications. In *Proceedings of the 17th ACM SIGSPATIAL international conference on advances in geographic information systems*, pages 246–255, 2009. 18

[99] Philippe Golle and Kurt Partridge. On the anonymity of home/work location pairs. In *International Conference on Pervasive Computing*, pages 390–397. Springer, 2009. 5

[100] Google. *Waze*, 2017. URL https://www.waze.com. 98

[101] Frank Gouineau, Tom Landry, and Thomas Triplet. Patchwork, a scalable density-grid clustering algorithm. In *Proceedings of the 31st Annual ACM Symposium on Applied Computing*, pages 824–831, 2016. 23

[102] Marco Gramaglia and Marco Fiore. Hiding mobile traffic fingerprints with glove. In *Proceedings of the 11th ACM Conference on Emerging Networking Experiments and Technologies*, pages 1–13, 2015. 27

[103] Match Group. Tinder, 2012. URL https://tinder.com/. 4

[104] Marco Gruteser and Dirk Grunwald. Anonymous usage of location-based services through spatial and temporal cloaking. In *Proceedings of the 1st international conference on Mobile systems, applications and services*, pages 31–42, 2003. 27

[105] Marco Gruteser and Baik Hoh. On the anonymity of periodic location samples. In *International Conference on Security in Pervasive Computing*, pages 179–192. Springer, 2005. 26

[106] Tianyu Gu, Brendan Dolan-Gavitt, and Siddharth Garg. Badnets: Identifying vulnerabilities in the machine learning model supply chain. *arXiv preprint arXiv:1708.06733*, 2017. 94

[107] Rachid Guerraoui, Sébastien Rouault, et al. The hidden vulnerability of distributed learning in byzantium. In *International Conference on Machine Learning*, pages 3521–3530. PMLR, 2018. 40

[108] Saikat Guha, Mudit Jain, and Venkata N Padmanabhan. Koi: A {Location-Privacy} platform for smartphone apps. In *9th USENIX Symposium on Networked Systems Design and Implementation (NSDI 12)*, pages 183–196, 2012. 30

[109] Xiaojie Guo, Zheli Liu, Jin Li, Jiqiang Gao, Boyu Hou, Changyu Dong, and Thar Baker. Verifl: Communication-efficient and fast verifiable aggregation for federated learning. *IEEE Transactions on Information Forensics and Security*, 16:1736–1751, 2020. 95

[110] Ruchika Gupta and Udai Pratap Rao. An exploration to location based service and its privacy preserving techniques: A survey. *Wireless Personal Communications*, 96(2):1973–2007, 2017. doi: 10.1007/s11277-017-4284-2. URL https://doi.org/10.1007/s11277-017-4284-2. 44

[111] Nicolas Haderer, Romain Rouvoy, Christophe Ribeiro, and Lionel Seinturier. Apisense: Crowd-sensing made easy. *ERCIM News*, 93:28–29, 2013. 4

[112] Andrew Hard, Kanishka Rao, Rajiv Mathews, Swaroop Ramaswamy, Françoise Beaufays, Sean Augenstein, Hubert Eichner, Chloé Kiddon, and Daniel Ramage. Federated learning for mobile keyboard prediction. *arXiv preprint arXiv:1811.03604*, 2018. 38, 70, 71

[113] István Hegedűs, Gábor Danner, and Márk Jelasity. Gossip learning as a decentralized alternative to federated learning. In *IFIP International Conference on Distributed Applications and Interoperable Systems*, pages 74–90. Springer, 2019. 138

[114] Janine S Hiller and Roberta S Russell. Privacy in Crises: The NIST Privacy Framework. *Journal of Contingencies and Crisis Management*, 25(1):31–38, 2017. 6, 68

[115] Briland Hitaj, Giuseppe Ateniese, and Fernando Perez-Cruz. Deep models under the gan: information leakage from collaborative deep learning. In *Proceedings of the 2017 ACM SIGSAC conference on computer and communications security*, pages 603–618, 2017. 41

[116] Sepp Hochreiter and Jürgen Schmidhuber. Long Short-Term Memory. *Neural Computation*, 9(8):1735–1780, 1997. 74

[117] Sameera Horawalavithana, Clayton Gandy, Juan Arroyo Flores, John Skvoretz, and Adriana Iamnitchi. Diversity, homophily and the risk of node re-identification in labeled social graphs. In Luca Maria Aiello, Chantal Cherifi, Hocine Cherifi, Renaud Lambiotte, Pietro Lió, and Luis Mateus Rocha, editors, *Complex Networks and Their Applications VII - Volume 2 Proceedings The 7th International Conference on Complex Networks and Their Applications COMPLEX NETWORKS 2018*, volume 813 of *Studies in Computational Intelligence*, pages 400–411. Springer, 2018. doi: 10.1007/978-3-030-05414-4_32. URL https://doi.org/10.1007/978-3-030-05414-4_32. 21

[118] Yan Huang, Zhipeng Cai, and Anu G. Bourgeois. Search locations safely and accurately: A location privacy protection algorithm with accurate service. *J. Network and Computer Applications*, 103:146–156, 2018. doi: 10.1016/j.jnca.2017.12.002. URL https://doi.org/10.1016/j.jnca.2017.12.002. 25, 45, 56

[119] Yan Huang, Zhipeng Cai, and Anu G Bourgeois. Search locations safely and accurately: A location privacy protection algorithm with accurate service. *Journal of Network and Computer Applications*, 103:146–156, 2018. 28, 98, 101, 115

[120] Yahoo! Inc. *Yahoo Weather*, 2013. URL https://mobile.yahoo.com/weather. 4

[121] Jason Jacobs. Runkeeper, 2008. URL https://runkeeper.com/cms/. 4

[122] Fengmei Jin, Wen Hua, Jiajie Xu, and Xiaofang Zhou. Moving object linking based on historical trace. In *2019 IEEE 35th International Conference on Data Engineering (ICDE)*, pages 1058–1069. IEEE, 2019. 23

[123] Jong Hee Kang, William Welbourne, Benjamin Stewart, and Gaetano Borriello. Extracting places from traces of locations. *ACM SIGMOBILE Mobile Computing and Communications Review*, 9(3):58–68, 2005. 18

[124] Ferhat Karakoç, Melek Önen, and Zeki Bilgin. Secure aggregation against malicious users. In *Proceedings of the 26th ACM Symposium on Access Control Models and Technologies*, pages 115–124, 2021. 41, 94

[125] Ilkcan Keles, Matthias Schubert, Peer Kröger, Simonas Šaltenis, and Christian S Jensen. Extracting visited points of interest from vehicle trajectories. In *Proceedings of the Fourth International ACM Workshop on Managing and Mining Enriched Geo-Spatial Data*, pages 1–6, 2017. 19

[126] Daniel Kelly, Barry Smyth, and Brian Caulfield. Uncovering measurements of social and demographic behavior from smartphone location data. *IEEE Transactions on Human-Machine Systems*, 43(2):188–198, 2013. 20

[127] Besma Khalfoun, Mohamed Maouche, Sonia Ben Mokhtar, and Sara Bouchenak. Mood: Mobility data privacy as orphan disease: Experimentation and deployment paper. In *Proceedings of the 20th International Middleware Conference, Middleware 2019, Davis, CA, USA, December 9-13, 2019*, pages 136–148. ACM, 2019. doi: 10.1145/3361525.3361542. URL https://doi.org/10.1145/3361525.3361542. 45, 98, 112

[128] Besma Khalfoun, Sonia Ben Mokhtar, Sara Bouchenak, and Vlad Nitu. EDEN: Enforcing Location Privacy through Re-identification Risk Assessment: A Federated Learning Approach. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies*, 5(2):1–25, 2021. 100

[129] Hidetoshi Kido, Yutaka Yanagisawa, and Tetsuji Satoh. An anonymous communication technique using dummies for location-based services. In *ICPS'05. Proceedings. International Conference on Pervasive Services, 2005.*, pages 88–97. IEEE, 2005. 28

[130] David G Kleinbaum, K Dietz, M Gail, Mitchel Klein, and Mitchell Klein. *Logistic Regression*. Springer, 2002. 78

[131] John Krumm. Inference attacks on location tracks. In Anthony LaMarca, Marc Langheinrich, and Khai N. Truong, editors, *Pervasive Computing, 5th*

*International Conference, PERVASIVE 2007, Toronto, Canada, May 13-16, 2007, Proceedings*, volume 4480 of *Lecture Notes in Computer Science*, pages 127–143. Springer, 2007. doi: 10.1007/978-3-540-72037-9_8. URL https://doi.org/10.1007/978-3-540-72037-9_8. 18, 104

[132] John Krumm. A Survey of Computational Location Privacy. *Personal and Ubiquitous Computing*, 13(6):391–399, 2009. 5

[133] John Krumm and Dany Rouhana. Placer: semantic place labels from diary data. In Friedemann Mattern, Silvia Santini, John F. Canny, Marc Langheinrich, and Jun Rekimoto, editors, *The 2013 ACM International Joint Conference on Pervasive and Ubiquitous Computing, UbiComp '13, Zurich, Switzerland, September 8-12, 2013*, pages 163–172. ACM, 2013. doi: 10.1145/2493432.2493504. URL https://doi.org/10.1145/2493432.2493504. 5, 19

[134] Vaibhav Kulkarni, Bertil Chapuis, and Benoît Garbinato. Privacy-preserving location-based services by using intel sgx. In *Proceedings of the First International Workshop on Human-Centered Sensing, Networking, and Systems*, pages 13–18, 2017. 31

[135] Juha K Laurila, Daniel Gatica-Perez, Imad Aad, Olivier Bornet, Trinh-Minh-Tri Do, Olivier Dousse, Julien Eberle, Markus Miettinen, et al. The Mobile Data Challenge: Big Data for Mobile Computing Research. Technical report, 2012. 17, 57, 69, 79, 112

[136] Anliang Li, Shuang Wang, Wenzhu Li, Shengnan Liu, and Siyuan Zhang. Predicting human mobility with federated learning. In *Proceedings of the 28th International Conference on Advances in Geographic Information Systems*, pages 441–444, 2020. 41

[137] Fa Li, Zhipeng Gui, Zhaoyu Zhang, Dehua Peng, Siyu Tian, Kunxiaojia Yuan, Yunzeng Sun, Huayi Wu, Jianya Gong, and Yichen Lei. A hierarchical temporal attention-based lstm encoder-decoder model for individual mobility prediction. *Neurocomputing*, 403:153–166, 2020. 20

[138] Huaxin Li, Haojin Zhu, Suguo Du, Xiaohui Liang, and Xuemin Shen. Privacy leakage of location sharing in mobile social networks: Attacks and defense. *IEEE Transactions on Dependable and Secure Computing*, 15(4):646–660, 2016. 106

[139] Huaxin Li, Haojin Zhu, Suguo Du, Xiaohui Liang, and Xuemin Sherman Shen. Privacy leakage of location sharing in mobile social networks: Attacks and defense. *IEEE Transactions on Dependable and Secure Computing*, 15(4): 646–660, 2018. 32

[140] Jie Li, Fanzi Zeng, Zhu Xiao, Hongbo Jiang, Zhirun Zheng, Wenping Liu, and Ju Ren. Drive2friends: Inferring social relationships from individual vehicle mobility data. *IEEE Internet of Things Journal*, 7(6):5116–5127, 2020. 20

[141] Jie Li, Fanzi Zeng, Zhu Xiao, Zhirun Zheng, Hongbo Jiang, and Zhetao Li. Social relationship inference over private vehicle mobility data. *IEEE Transactions on Vehicular Technology*, 70(6):5221–5233, 2021. 20

[142] Tian Li, Anit Kumar Sahu, Ameet Talwalkar, and Virginia Smith. Federated learning: Challenges, methods, and future directions. *IEEE Signal Process. Mag.*, 37(3):50–60, 2020. doi: 10.1109/MSP.2020.2975749. URL https://doi.org/10.1109/MSP.2020.2975749. 70

[143] Zhaorui Li, Zhicong Huang, Chaochao Chen, and Cheng Hong. Quantification of the leakage in federated learning. *arXiv preprint arXiv:1910.05467*, 2019. 41, 95

[144] Yang Liu, Yan Kang, Chaoping Xing, Tianjian Chen, and Qiang Yang. A secure federated transfer learning framework. *IEEE Intelligent Systems*, 35(4): 70–82, 2020. 40

[145] Zheng Liu, Yuanyuan Qiao, Siyan Tao, Wenhui Lin, and Jie Yang. Analyzing human mobility and social relationships from cellular network data. In *2017 13th International Conference on Network and Service Management (CNSM)*, pages 1–6. IEEE, 2017. 20

[146] Xin Lu, Erik Wetter, Nita Bharti, Andrew J Tatem, and Linus Bengtsson. Approaching the limit of predictability in human mobility. *Scientific reports*, 3 (1):1–9, 2013. 128

[147] Lingjuan Lyu, Jiangshan Yu, Karthik Nandakumar, Yitong Li, Xingjun Ma, and Jiong Jin. Towards fair and decentralized privacy-preserving deep learning with blockchain. *arXiv preprint arXiv:1906.01167*, pages 1–13, 2019. 39

[148] Zhenliang Ma and Pengfei Zhang. Individual mobility prediction review: Data, problem, method and application. *Multimodal Transportation*, 1(1):100002, 2022. 20

[149] Ashwin Machanavajjhala, Daniel Kifer, Johannes Gehrke, and Muthuramakrishnan Venkitasubramaniam. l-diversity: Privacy beyond k-anonymity. *ACM Transactions on Knowledge Discovery from Data (TKDD)*, 1(1):3–es, 2007. 33

[150] Nicolas Maisonneuve, Matthias Stevens, Maria E. Niessen, and Luc Steels. Noisetube: Measuring and mapping noise pollution with mobile phones. In *Information Technologies in Environmental Engineering, Proceedings of the 4th International ICSC Symposium, ITEE 2009, Thessaloniki, Greece, May 28-29, 2009*, pages 215–228, 2009. doi: 10.1007/978-3-540-88351-7_16. URL https://doi.org/10.1007/978-3-540-88351-7_16. 63

[151] Nicolas Maisonneuve, Matthias Stevens, Maria E Niessen, and Luc Steels. Noisetube: Measuring and mapping noise pollution with mobile phones. In *Information technologies in environmental engineering*, pages 215–228. Springer, 2009. 4

[152] Mohamed Maouche, Sonia Ben Mokhtar, and Sara Bouchenak. Ap-attack: A novel user re-identification attack on mobility datasets. In Tao Gu, Ramamohanarao Kotagiri, and Huai Liu, editors, *Proceedings of the 14th EAI International Conference on Mobile and Ubiquitous Systems: Computing, Networking and Services, Melbourne, Australia, November 7-10, 2017*, pages 48–57. ACM, 2017. doi: 10.1145/3144457.3144494. URL https://doi.org/10.1145/3144457.3144494. 5, 21, 22, 31, 45, 46, 56, 57, 69, 90, 98, 99, 105, 106, 115

[153] Mohamed Maouche, Sonia Ben Mokhtar, and Sara Bouchenak. Hmc. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies*, 2(3):1–25, 2018. 30, 31, 34, 45, 46, 54, 56, 57, 62, 112

[154] Sergio Mascetti, Claudio Bettini, X Sean Wang, Dario Freni, and Sushil Jajodia. Providenthider: An algorithm to preserve historical k-anonymity in lbs. In *2009 Tenth International Conference on Mobile Data Management: Systems, Services and Middleware*, pages 172–181. IEEE, 2009. 36

[155] Sergio Mascetti, Dario Freni, Claudio Bettini, X Sean Wang, and Sushil Jajodia. Privacy in geo-social networks: proximity notification with untrusted service providers and curious buddies. *The VLDB journal*, 20(4):541–566, 2011. 30

[156] Amirreza Masoumzadeh and James Joshi. An alternative approach to k-anonymity for location-based services. *Procedia Computer Science*, 5:522–530, 2011. 36

[157] Clément Massart and François-Xavier Standaert. Revisiting location privacy from a side-channel analysis viewpoint (extended version). *IACR Cryptol. ePrint Arch.*, page 467, 2019. URL https://eprint.iacr.org/2019/467. 23

[158] Wesley Mathew, Ruben Raposo, and Bruno Martins. Predicting future locations with hidden markov models. In *Proceedings of the 2012 ACM conference on ubiquitous computing*, pages 911–918, 2012. 20

[159] Brendan McMahan, Eider Moore, Daniel Ramage, Seth Hampson, and Blaise Aguera y Arcas. Communication-efficient learning of deep networks from decentralized data. In *Artificial Intelligence and Statistics*, pages 1273–1282. PMLR, 2017. 38, 70, 99

[160] Scott Menard. *Applied logistic regression analysis*, volume 106. Sage, 2002. 105

[161] Aghiles Ait Messaoud, Sonia Ben Mokhtar, Vlad Nitu, and Valerio Schiavoni. Gradsec: a tee-based scheme against federated learning inference attacks. In *Proceedings of the First Workshop on Systems Challenges in Reliable and Secure Federated Learning*, pages 10–12, 2021. 42, 94

[162] Kristopher Micinski, Philip Phelps, and Jeffrey S Foster. An empirical study of location truncation on android. *Weather*, 2:21, 2013. 29

[163] Tomas Mikolov, Kai Chen, Greg Corrado, and Jeffrey Dean. Efficient Estimation of Word Representations in Vector Space. *arXiv preprint arXiv:1301.3781*, 2013. 75

[164] Fan Mo and Hamed Haddadi. Efficient and private federated learning using tee. In *EuroSys*, 2019. 131, 136

[165] Fan Mo, Ali Shahin Shamsabadi, Kleomenis Katevas, Soteris Demetriou, Ilias Leontiadis, Andrea Cavallaro, and Hamed Haddadi. Darknetz: towards model

privacy at the edge using trusted execution environments. In *Proceedings of the 18th International Conference on Mobile Systems, Applications, and Services*, pages 161–174, 2020. 42, 94

[166] Fan Mo, Hamed Haddadi, Kleomenis Katevas, Eduard Marin, Diego Perino, and Nicolas Kourtellis. Ppfl: privacy-preserving federated learning with trusted execution environments. In *Proceedings of the 19th Annual International Conference on Mobile Systems, Applications, and Services*, pages 94–108, 2021. 95

[167] Prashanth Mohan, Venkata N Padmanabhan, and Ramachandran Ramjee. Nericell: rich monitoring of road and traffic conditions using mobile smartphones. In *Proceedings of the 6th ACM conference on Embedded network sensor systems*, pages 323–336, 2008. 4

[168] Mohamed F Mokbel, Chi-Yin Chow, and Walid G Aref. The new casper: Query processing for location services without compromising privacy. In *VLDB*, volume 6, pages 763–774, 2006. 27

[169] Sonia Ben Mokhtar, Antoine Boutet, Louafi Bouzouina, Patrick Bonnel, Olivier Brette, Lionel Brunie, Mathieu Cunche, Stephane D'Alu, Vincent Primault, Patrice Raveneau, et al. Priva'mov: Analysing Human Mobility Through Multi-Sensor Datasets. In *NetMob 2017*, 2017. 17, 57, 69, 79, 101, 112

[170] Raul Montoliu, Jan Blom, and Daniel Gatica-Perez. Discovering places of interest in everyday life from smartphone data. *Multimedia tools and applications*, 62(1):179–207, 2013. 18

[171] Ahmed Moustafa, Muhammad Asad, Saima Shaukat, and Alexander Norta. Ppcsa: Partial participation-based compressed and secure aggregation in federated learning. In *International Conference on Advanced Information Networking and Applications*, pages 345–357. Springer, 2021. 41, 95

[172] Farid M Naini, Jayakrishnan Unnikrishnan, Patrick Thiran, and Martin Vetterli. Where you are is who you are: User identification by matching statistics. *IEEE Transactions on Information Forensics and Security*, 11(2):358–372, 2015. 22, 106

[173] Arvind Narayanan and Vitaly Shmatikov. Robust de-anonymization of large sparse datasets. In *2008 IEEE Symposium on Security and Privacy (S&P 2008), 18-21 May 2008, Oakland, California, USA*, pages 111–125. IEEE Computer Society, 2008. doi: 10.1109/SP.2008.33. URL https://doi.org/10.1109/SP.2008.33. 21

[174] Arvind Narayanan and Vitaly Shmatikov. De-anonymizing social networks. In *30th IEEE Symposium on Security and Privacy (S&P 2009), 17-20 May 2009, Oakland, California, USA*, pages 173–187. IEEE Computer Society, 2009. doi: 10.1109/SP.2009.22. URL https://doi.org/10.1109/SP.2009.22. 21

[175] Francesca Naretto, Roberto Pellungrini, Anna Monreale, Franco Maria Nardini, and Mirco Musolesi. Predicting and Explaining Privacy Risk Exposure in Mobility Data. In *International Conference on Discovery Science*, pages 403–418. Springer, 2020. 35

[176] Milad Nasr, Reza Shokri, and Amir Houmansadr. Comprehensive privacy analysis of deep learning: Passive and active white-box inference attacks against centralized and federated learning. In *2019 IEEE Symposium on Security and Privacy (SP)*, pages 739–753. IEEE, 2019. 41, 95

[177] Thien Duc Nguyen, Phillip Rieger, Hossein Yalame, Helen Möllering, Hossein Fereidooni, Samuel Marchal, Markus Miettinen, Azalia Mirhoseini, Ahmad-Reza Sadeghi, Thomas Schneider, et al. *arXiv preprint arXiv:2101.02281*, 2021. 94

[178] Niantic. Pokemon Go, 2017. URL http://www.pokemongo.com. 4

[179] Gunarto Sindoro Njoo, Min-Chia Kao, Kuo-Wei Hsu, and Wen-Chih Peng. Exploring check-in data to infer social ties in location based social networks. In *Pacific-Asia Conference on Knowledge Discovery and Data Mining*, pages 460–471. Springer, 2017. 20

[180] Sankar K. Pal and Sushmita Mitra. Multilayer perceptron, fuzzy sets, and classification. *IEEE Trans. Neural Networks*, 3(5):683–697, 1992. doi: 10.1109/72.159058. URL https://doi.org/10.1109/72.159058. 105

[181] Balaji Palanisamy and Ling Liu. Mobimix: Protecting location privacy with mix-zones over road networks. In *2011 IEEE 27th International conference on data engineering*, pages 494–505. IEEE, 2011. 26

[182] European Parliament and council. *European Union genral data protection regulation*, 2016. URL https://eur-lex.europa.eu/legal-content/EN/TXT/?qid=1465452422595&uri=CELEX:32016R0679. 6

[183] European Parliament and council. *European Union genral data protection regulation*, 2016. URL https://gdpr-info.eu/issues/privacy-by-design/. 6

[184] Dario Pasquini, Danilo Francati, and Giuseppe Ateniese. Eluding secure aggregation in federated learning via model inconsistency. *arXiv preprint arXiv:2111.07380*, 2021. 41, 95

[185] Sai Teja Peddinti and Nitesh Saxena. On the limitations of query obfuscation techniques for location privacy. In *Proceedings of the 13th international conference on Ubiquitous computing*, pages 187–196, 2011. 28

[186] Roberto Pellungrini, Luca Pappalardo, Francesca Pratesi, and Anna Monreale. A Data Mining Approach to Assess Privacy Risk in Human Mobility Data. *ACM Transactions on Intelligent Systems and Technology (TIST)*, 9(3):1–27, 2017. 36

[187] Roberto Pellungrini, Luca Pappalardo, Francesca Pratesi, and Anna Monreale. Fast Estimation of Privacy Risk in Human Mobility Data. In *International Conference on Computer Safety, Reliability, and Security*, pages 415–426. Springer, 2017. 36

[188] Huy Pham, Cyrus Shahabi, and Yan Liu. Ebm: an entropy-based model to infer social strength from spatiotemporal data. In *Proceedings of the 2013 ACM SIGMOD international conference on management of data*, pages 265–276, 2013. 20

[189] Aniket Pingley, Wei Yu, Nan Zhang, Xinwen Fu, and Wei Zhao. Cap: A context-aware privacy protection system for location-based services. In *2009 29th IEEE International Conference on Distributed Computing Systems*, pages 49–57. IEEE, 2009. 29

[190] Michal Piorkowski, Natasa Sarafijanovic-Djukic, and Matthias Grossglauser. Crawdad data set epfl/mobility (v. 2009-02-24), 2009. 57, 117, 121

[191] Layla Pournajaf, Daniel A Garcia-Ulloa, Li Xiong, and Vaidy Sunderam. Participant privacy in mobile crowd sensing task management: A survey of methods and challenges. *ACM Sigmod Record*, 44(4):23–34, 2016. 104

[192] Francesca Pratesi, Anna Monreale, Roberto Trasarti, Fosca Giannotti, Dino Pedreschi, and Tadashi Yanagihara. Prudence: a system for assessing privacy risk vs utility in data sharing ecosystems. 2018. 35, 98

[193] Vincent Primault, Sonia Ben Mokhtar, Cédric Lauradoux, and Lionel Brunie. Differentially private location privacy in practice. *arXiv preprint arXiv:1410.7744*, 2014. 5, 45, 70, 79, 98, 99, 105, 112, 115

[194] Vincent Primault, Sonia Ben Mokhtar, Cédric Lauradoux, and Lionel Brunie. Time distortion anonymization for the publication of mobility data with high utility. In *2015 IEEE TrustCom/BigDataSE/ISPA, Helsinki, Finland, August 20-22, 2015, Volume 1*, pages 539–546. IEEE, 2015. doi: 10.1109/Trustcom.2015.417. URL https://doi.org/10.1109/Trustcom.2015.417. 22

[195] Vincent Primault, Sonia Ben Mokhtar, Cédric Lauradoux, and Lionel Brunie. Time distortion anonymization for the publication of mobility data with high utility. *CoRR*, abs/1507.00443, 2015. URL http://arxiv.org/abs/1507.00443. 25, 30, 31, 33

[196] Vincent Primault, Sonia Ben Mokhtar, Cédric Lauradoux, and Lionel Brunie. Time distortion anonymization for the publication of mobility data with high utility. In *2015 IEEE Trustcom/BigDataSE/ISPA*, volume 1, pages 539–546. IEEE, 2015. 34, 98, 101, 112, 113, 115

[197] Vincent Primault, Antoine Boutet, Sonia Ben Mokhtar, and Lionel Brunie. Adaptive location privacy with ALP. In *35th IEEE Symposium on Reliable Distributed Systems, SRDS 2016, Budapest, Hungary, September 26-29, 2016*, pages 269–278. IEEE Computer Society, 2016. doi: 10.1109/SRDS.2016.044. URL https://doi.org/10.1109/SRDS.2016.044. 56

[198] Vincent Primault, Antoine Boutet, Sonia Ben Mokhtar, and Lionel Brunie. Adaptive location privacy with alp. In *2016 IEEE 35th Symposium on Reliable Distributed Systems (SRDS)*, pages 269–278. IEEE, 2016. 32, 34, 101, 108

[199] Vincent Primault, Antoine Boutet, Sonia Ben Mokhtar, and Lionel Brunie. The Long Road to Computational Location Privacy: A Survey. *Communications Surveys and Tutorials, IEEE Communications Society*, page 1, 2018. doi: 10.1109/COMST.2018.2873950. URL https://hal.archives-ouvertes.fr/hal-01890014. 18, 21, 24, 32, 34, 44, 98, 113

[200] Vincent Primault, Mohamed Maouche, Antoine Boutet, Sonia Ben Mokhtar, Sara Bouchenak, and Lionel Brunie. ACCIO: how to make location privacy experimentation open and easy. In *38th IEEE International Conference on Distributed Computing Systems, ICDCS 2018, Vienna, Austria, July 2-6, 2018*, pages 896–906. IEEE Computer Society, 2018. doi: 10.1109/ICDCS.2018.00091. URL https://doi.org/10.1109/ICDCS.2018.00091. 55

[201] Apostolos Pyrgelis, Carmela Troncoso, and Emiliano De Cristofaro. Knock knock, who's there? membership inference on aggregate location data. *In NDSS*, 2018. 41

[202] Jean Louis Raisaro, Florian Tramèr, Zhanglong Ji, Diyue Bu, Yongan Zhao, W. Knox Carey, David D. Lloyd, Heidi Sofia, Dixie Baker, Paul Flicek, Suyash S. Shringarpure, Carlos D. Bustamante, Shuang Wang, Xiaoqian Jiang, Lucila Ohno-Machado, Haixu Tang, XiaoFeng Wang, and Jean-Pierre Hubaux. Addressing beacon re-identification attacks: quantification and mitigation of privacy risks. *J. Am. Medical Informatics Assoc.*, 24(4):799–805, 2017. doi: 10.1093/jamia/ocw167. URL https://doi.org/10.1093/jamia/ocw167. 21

[203] Daniele Riboni and Claudio Bettini. Differentially-private release of check-in data for venue recommendation. In *IEEE International Conference on Pervasive Computing and Communications, PerCom 2014, Budapest, Hungary, March 24-28, 2014*, pages 190–198. IEEE Computer Society, 2014. doi: 10.1109/PerCom.2014.6813960. URL https://doi.org/10.1109/PerCom.2014.6813960. 34

[204] John E Roemer. Equality of opportunity: A progress report. *Social Choice and Welfare*, 19(2):455–471, 2002. 128

[205] Nick Mathewson Roger Dingledine. *The TOR Project*, 2006. URL https://www.torproject.org. 100, 104

[206] Marco Romanelli, Catuscia Palamidessi, and Konstantinos Chatzikokolakis.

Generating optimal privacy-protection mechanisms via machine learning. *arXiv preprint arXiv:1904.01059*, 2019. 30, 98

[207] Luca Rossi and Mirco Musolesi. It's the way you check-in: Identifying users in location-based social networks. In *Proceedings of the second ACM conference on Online social networks*, pages 215–226, 2014. 22

[208] Luca Rossi, James Walker, and Mirco Musolesi. Spatio-Temporal Techniques for User Identification by Means of GPS Mobility Data. *EPJ Data Science*, 4 (1):11, 2015. 36

[209] Adam Sadilek and John Krumm. Far out: Predicting long-term human mobility. In Jörg Hoffmann and Bart Selman, editors, *Proceedings of the Twenty-Sixth AAAI Conference on Artificial Intelligence, July 22-26, 2012, Toronto, Ontario, Canada*. AAAI Press, 2012. URL http://www.aaai.org/ocs/index.php/AAAI/AAAI12/paper/view/4845. 20

[210] Aniruddha Saha, Akshayvarun Subramanya, and Hamed Pirsiavash. Hidden trigger backdoor attacks. In *Proceedings of the AAAI conference on artificial intelligence*, volume 34, pages 11957–11965, 2020. 41

[211] Julián Salas, David Megías, and Vicenç Torra. Swapmob: Swapping trajectories for mobility anonymization. In *International Conference on Privacy in Statistical Databases*, pages 331–346. Springer, 2018. 26

[212] Pierangela Samarati. Protecting respondents identities in microdata release. *IEEE transactions on Knowledge and Data Engineering*, 13(6):1010–1027, 2001. 98

[213] James Scott, Richard Gass, Jon Crowcroft, Pan Hui, Christophe Diot, and Augustin Chaintreau. Crawdad dataset cambridge/haggle (v. 2006-09-15). *CRAWDAD wireless network data archive*, 2006. 17, 24

[214] Mohamed Seif, Ravi Tandon, and Ming Li. Wireless federated learning with local differential privacy. In *2020 IEEE International Symposium on Information Theory (ISIT)*, pages 2604–2609. IEEE, 2020. 41, 95

[215] Ali Shafahi, W Ronny Huang, Mahyar Najibi, Octavian Suciu, Christoph Studer, Tudor Dumitras, and Tom Goldstein. Poison frogs! targeted clean-

label poisoning attacks on neural networks. *Advances in neural information processing systems*, 31, 2018. 41, 94

[216] Zhangqing Shan, Weiwei Sun, and Baihua Zheng. Extract human mobility patterns powered by city semantic diagram. *IEEE Transactions on Knowledge and Data Engineering*, 2020. 18

[217] Pravin Shankar, Vinod Ganapathy, and Liviu Iftode. Privately querying location-based services with sybilquery. In *Proceedings of the 11th international conference on Ubiquitous computing*, pages 31–40, 2009. 28

[218] Shiqi Shen, Shruti Tople, and Prateek Saxena. Auror: Defending against poisoning attacks in collaborative deep learning systems. In *Proceedings of the 32nd Annual Conference on Computer Security Applications*, pages 508–519, 2016. 41, 42, 94

[219] Reza Shokri, George Theodorakopoulos, Jean-Yves Le Boudec, and Jean-Pierre Hubaux. Quantifying location privacy. In *32nd IEEE Symposium on Security and Privacy, S&P 2011, 22-25 May 2011, Berkeley, California, USA*, pages 247–262. IEEE Computer Society, 2011. doi: 10.1109/SP.2011.18. URL https://doi.org/10.1109/SP.2011.18. 18

[220] Reza Shokri, George Theodorakopoulos, Jean-Yves Le Boudec, and Jean-Pierre Hubaux. Quantifying location privacy. In *2011 IEEE symposium on security and privacy*, pages 247–262. IEEE, 2011. 18

[221] Yi Song, Daniel Dahlmeier, and Stephane Bressan. Not So Unique in the Crowd: A Simple and Effective Algorithm for Anonymizing Location Data. In *The First International Workshop on Privacy-Preserving IR: When Information Retrieval Meets Privacy and Security (PIR 2014)*, 2014. 36, 68

[222] Mudhakar Srivatsa and Michael Hicks. Deanonymizing mobility traces: using social network as a side-channel. In Ting Yu, George Danezis, and Virgil D. Gligor, editors, *the ACM Conference on Computer and Communications Security, CCS'12, Raleigh, NC, USA, October 16-18, 2012*, pages 628–637. ACM, 2012. doi: 10.1145/2382196.2382262. URL https://doi.org/10.1145/2382196.2382262. 21, 23

[223] Ziteng Sun, Peter Kairouz, Ananda Theertha Suresh, and H Brendan McMahan. Can you really backdoor federated learning? *arXiv preprint arXiv:1911.07963*, 2019. 136

[224] Latanya Sweeney. k-anonymity: A model for protecting privacy. *International Journal of Uncertainty, Fuzziness and Knowledge-Based Systems*, 10(05):557–570, 2002. 27, 33, 36

[225] Stuart A. Thompson and Charlie Warzel. *How to Track President Trump*, 2019. URL https://www.nytimes.com/interactive/2019/12/20/opinion/location-data-national-security.html. 5

[226] Vale Tolpegin, Stacey Truex, Mehmet Emre Gursoy, and Ling Liu. Data Poisoning Attacks Against Federated Learning Systems. In *European Symposium on Research in Computer Security*, pages 480–501. Springer, 2020. 94

[227] Aleksei Triastcyn and Boi Faltings. Federated learning with bayesian differential privacy. In *2019 IEEE International Conference on Big Data (Big Data)*, pages 2587–2596. IEEE, 2019. 41, 95

[228] Zhen Tu, Runtong Li, Yong Li, Gang Wang, Di Wu, Pan Hui, Li Su, and Depeng Jin. Your Apps Give You Away: Distinguishing Mobile Users by Their App Usage Fingerprints. *Proc. ACM Interact. Mob. Wearable Ubiquitous Technol.*, 2(3):138:1–138:23, 2018. doi: 10.1145/3264948. URL https://doi.org/10.1145/3264948. 36

[229] Paul Vanhaesebrouck, Aurélien Bellet, and Marc Tommasi. Decentralized collaborative learning of personalized models over networks. In *Artificial Intelligence and Statistics*, pages 509–517. PMLR, 2017. 39

[230] Antoine Vastel, Pierre Laperdrix, Walter Rudametkin, and Romain Rouvoy. FP-STALKER: tracking browser fingerprint evolutions. In *2018 IEEE Symposium on Security and Privacy, SP 2018, Proceedings, 21-23 May 2018, San Francisco, California, USA*, pages 728–741. IEEE Computer Society, 2018. doi: 10.1109/SP.2018.00008. URL https://doi.org/10.1109/SP.2018.00008. 21

[231] Raj Kiriti Velicheti, Derek Xia, and Oluwasanmi Koyejo. Secure byzantine-robust distributed learning via clustering. *arXiv preprint arXiv:2110.02940*, 2021. 94

[232] Isabel Wagner and David Eckhoff. Technical privacy metrics: a systematic survey. *ACM Computing Surveys (CSUR)*, 51(3):1–38, 2018. 32, 33

[233] Dashun Wang, Dino Pedreschi, Chaoming Song, Fosca Giannotti, and Albert-Laszlo Barabasi. Human mobility, social ties, and link prediction. In *Proceedings of the 17th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 1100–1108. Acm, 2011. 5, 19

[234] Huandong Wang, Chen Gao, Yong Li, Zhi-Li Zhang, and Depeng Jin. From Fingerprint to Footprint: Revealing Physical World Privacy Leakage by Cyberspace Cookie Logs. In *Proceedings of the 2017 ACM on Conference on Information and Knowledge Management*, pages 1209–1218, 2017. 5, 36

[235] Huandong Wang, Chen Gao, Yong Li, Gang Wang, Depeng Jin, and Jingbo Sun. De-anonymization of mobility trajectories: Dissecting the gaps between theory and practice. In *The 25th Annual Network & Distributed System Security Symposium (NDSS'18)*, 2018. 26

[236] Jiangtao Wang, Yasha Wang, Daqing Zhang, and Sumi Helal. Energy saving techniques in mobile crowd sensing: Current state and future opportunities. *IEEE Communications Magazine*, 56(5):164–169, 2018. 112

[237] Yan Wang, Ali Yalcin, and Carla VandeWeerd. An entropy-based approach to the study of human mobility and behavior in private homes. *PLoS one*, 15(12): e0243503, 2020. 128

[238] Marius Wernke, Pavel Skvortsov, Frank Dürr, and Kurt Rothermel. A classification of location privacy attacks and approaches. *Pers. Ubiquitous Comput.*, 18(1):163–175, 2014. doi: 10.1007/s00779-012-0633-z. URL https://doi.org/10.1007/s00779-012-0633-z. 18

[239] Marius Wernke, Pavel Skvortsov, Frank Dürr, and Kurt Rothermel. A classification of location privacy attacks and approaches. *Personal and ubiquitous computing*, 18(1):163–175, 2014. 5

[240] Jeff Williams. Owasp risk rating methodology. *OWASP. Available online: https://owasp. org/www-community/OWASP_Risk_ Rating_Methodology (accessed on 11 January 2021)*, 2020. 35

[241] Guowen Xu, Hongwei Li, Sen Liu, Kan Yang, and Xiaodong Lin. Verifynet: Secure and verifiable federated learning. *IEEE Transactions on Information Forensics and Security*, 15:911–926, 2019. 95

[242] Toby Xu and Ying Cai. Exploring historical location data for anonymity preservation in location-based services. In *IEEE INFOCOM 2008-The 27th Conference on Computer Communications*, pages 547–555. IEEE, 2008. 36

[243] Zhikai Xu, Hongli Zhang, and Xiangzhan Yu. Multiple mix-zones deployment for continuous location privacy protection. In *2016 IEEE Trustcom/BigDataSE/ISPA*, pages 760–766. IEEE, 2016. 26

[244] Mingqiang Xue, Panos Kalnis, and Hung Keng Pung. Location diversity: Enhanced privacy protection in location based services. In *International Symposium on Location-and Context-Awareness*, pages 70–87. Springer, 2009. 33

[245] Manxiang Yang, Baopeng Ye, Yuling Chen, Tao Li, Yixian Yang, Xiaobin Qian, and Xiaomei Yu. A trusted de-swinging k-anonymity scheme for location privacy protection. *Journal of Cloud Computing*, 11(1):1–15, 2022. 37

[246] Timothy Yang, Galen Andrew, Hubert Eichner, Haicheng Sun, Wei Li, Nicholas Kong, Daniel Ramage, and Françoise Beaufays. Applied federated learning: Improving google keyboard query suggestions. *arXiv preprint arXiv:1812.02903*, 2018. 38

[247] Ting-Fang Yen, Yinglian Xie, Fang Yu, Roger Peng Yu, and Martín Abadi. Host Fingerprinting and Tracking on the Web: Privacy and Security Implications. In *The 19th Annual Network and Distributed System Security Symposium (NDSS 2012)*, San Diego, California, USA, 2012. URL https://www.ndss-symposium.org/ndss2012/host-fingerprinting-and-tracking-web-privacy-and-security-implications. 36

[248] Sandy L Zabell. Fingerprint Evidence. *JL & Pol'y*, 13:143, 2005. 36

[249] Polixeni Zacharouli, Aris Gkoulalas-Divanis, and Vassilios S Verykios. A k-anonymity model for spatio-temporal data. In *2007 IEEE 23rd International Conference on Data Engineering Workshop*, pages 555–564. IEEE, 2007. 36

[250] Hui Zang and Jean Bolot. Anonymization of Location Data Does Not Work: A Large-Scale Measurement Study. In Parmesh Ramanathan, Thyaga Nandagopal, and Brian Neil Levine, editors, *Proceedings of the 17th Annual International Conference on Mobile Computing and Networking, MOBICOM 2011, Las Vegas, Nevada, USA, September 19-23, 2011*, pages 145–156. ACM, 2011. doi: 10.1145/2030613.2030630. URL https://doi.org/10.1145/2030613.2030630. 36

[251] Hongtao Zhang and Lingcheng Dai. Mobility prediction: A survey on state-of-the-art schemes and future applications. *IEEE access*, 7:802–822, 2018. 5, 20

[252] Jingwen Zhang, Jiale Zhang, Junjun Chen, and Shui Yu. GAN Enhanced Membership Inference: A Passive Local Attack in Federated Learning. In *ICC 2020-2020 IEEE International Conference on Communications (ICC)*, pages 1–6. IEEE, 2020. 94

[253] Xianglong Zhang, Anmin Fu, Huaqun Wang, Chunyi Zhou, and Zhenzhu Chen. A privacy-preserving and verifiable federated learning scheme. In *ICC 2020-2020 IEEE International Conference on Communications (ICC)*, pages 1–6. IEEE, 2020. 95

[254] Xinyang Zhang, Shouling Ji, Hui Wang, and Ting Wang. Private, yet practical, multiparty deep learning. In *2017 IEEE 37th International Conference on Distributed Computing Systems (ICDCS)*, pages 1442–1452. IEEE, 2017. 41, 71, 72

[255] Zhuosheng Zhang, Jiarui Li, Shucheng Yu, and Christian Makaya. Safelearning: Enable backdoor detectability in federated learning with secure aggregation. *arXiv preprint arXiv:2102.02402*, 2021. 41, 94

[256] Lingchen Zhao, Jianlin Jiang, Bo Feng, Qian Wang, Chao Shen, and Qi Li. Sear: Secure and efficient aggregation for byzantine-robust federated learning. *IEEE Transactions on Dependable and Secure Computing*, 2021. 42, 94, 95

[257] Yu Zheng, Xing Xie, Wei-Ying Ma, et al. Geolife: A collaborative social networking service among user, location and trajectory. *IEEE Data Eng. Bull.*, 33(2):32–39, 2010. 17, 57, 69, 79, 112

[258] Bin Zhou, Jian Pei, and WoShun Luk. A brief survey on anonymization techniques for privacy preserving publishing of social network data. *ACM Sigkdd Explorations Newsletter*, 10(2):12–22, 2008. 21

[259] Changqing Zhou, Dan Frankowski, Pamela Ludford, Shashi Shekhar, and Loren Terveen. Discovering personal gazetteers: an interactive clustering approach. In *Proceedings of the 12th annual ACM international workshop on Geographic information systems*, pages 266–273, 2004. 5, 56, 115

[260] Changqing Zhou, Dan Frankowski, Pamela J. Ludford, Shashi Shekhar, and Loren G. Terveen. Discovering personal gazetteers: an interactive clustering approach. In Dieter Pfoser, Isabel F. Cruz, and Marc Ronthaler, editors, *12th ACM International Workshop on Geographic Information Systems, ACM-GIS 2004, November 12-13, 2004, Washington, DC, USA, Proceedings*, pages 266–273. ACM, 2004. doi: 10.1145/1032222.1032261. URL https://doi.org/10.1145/1032222.1032261. 18

[261] Ligeng Zhu, Zhijian Liu, and Song Han. Deep leakage from gradients. *Advances in neural information processing systems*, 32, 2019. 41

FOLIO ADMINISTRATIF

THESE DE L'UNIVERSITE DE LYON OPEREE AU SEIN DE L'INSA LYON

NOM : KHALFOUN DATE de SOUTENANCE : 20/10/2022

Prénoms : Besma

TITRE : Privacy Preserving location based services: from centralized to federated approaches

NATURE : Doctorat Numéro d'ordre : 2022ISAL0089

Ecole doctorale : InfoMaths (ED 5120)

Spécialité :Informatique

RESUME :
De nos jours, la prolifération des appareils mobiles embarquant de multiples capteurs et la croissance rapide des technologies de communication et de traitement de la donnée ont contribué à l'émergence d'une grande variété de services en ligne, dont les services basés sur la localisation. Ces services facilitent la vie quotidienne des utilisateurs en leur offrant des informations personnalisées et customisées sur leur environnement en fonction de leur localisation. Tout en reconnaissant qu'il est indéniable que ces services sont devenus incontournables et indispensables à notre société actuelle et surtout future, il y a lieu de souligner et d'appréhender les risques et les dangers quant à la vie privée des utilisateurs. En effet, de grandes quantités de données de mobilité sont collectées, stockées et traitées par des fournisseurs de services ou des tiers, sans forcément respecter le cadre consenti par les utilisateurs. Par conséquent, la vie privée de ces derniers est menacée et donc plusieurs informations sensibles telles que l'identité de l'utilisateur, son lieu de domicile ou de travail ou même ses croyances religieuses ou son état de santé peuvent être inférées de ces données. Dans ce contexte, il devient urgent de concevoir des mécanismes de protection qui permettent aux utilisateurs d'accéder en toute sécurité aux services basés sur la localisation sans la crainte de dévoiler leur intimité. Pour relever ce défi, de nombreux efforts visent à développer des mécanismes de protection appelés « Location Privacy Protection Mechanisms (LPPM) ». Ces efforts ne sont pas seulement motivés par la communauté scientifique mais sont de plus en plus imposés par les autorités et les pouvoirs publics en établissant de nouvelles règles et lois pour recadrer la collecte, le stockage et la manipulation de ces données. Dans ce sens, l'évaluation des risques liés à la confidentialité de la mobilité appelée « Location Privacy Risk Assessment (LPRA) » est définie afin de sensibiliser les utilisateurs aux risques engendrés par le partage de leurs données de mobilité. Dans le cas de notre étude, cette notion se traduit par l'évaluation du risque de ré-identification, c'est-à-dire le risque de réassocier une donnée de mobilité anonyme à son utilisateur d'origine. Dans ce cadre, nous proposons tout d'abord MOOD, un système de protection centralisé centré sur l'utilisateur qui a pour but de protéger les données de mobilité de tous les utilisateurs et, en particulier, les utilisateurs orphelins qui ne sont protégés par aucun LPPM individuel. MOOD utilise la composition de plusieurs LPPM et intègre l'évaluation du risque de ré-identification avant de publier les données protégées. Cependant, il requiert un « serveur proxy de confiance » pour procéder à la protection et à l'évaluation du risque de ré-identification. Bien que les méthodes de protection actuelles tendent à éliminer ce serveur proxy de confiance, l'évaluation du risque d'atteinte à la vie privée a toujours besoin de centraliser les données de mobilité. Pour cette raison, nous proposons SAFER, une nouvelle mesure d'évaluation du risque de confidentialité, développée du côté utilisateur pour estimer le risque de confidentialité en utilisant l'unicité des données de mobilité appelée « uniqueness ». SAFER suit une approche basée sur l'apprentissage fédéré pour construire une connaissance globale sans avoir accès aux données brutes des utilisateurs de façon centralisée. Enfin, nous proposons EDEN, un système de protection des données de mobilité, développé du côté utilisateur. EDEN sélectionne automatiquement le meilleur LPPM et sa configuration correspondante sans envoyer les données de mobilité brutes en dehors du dispositif de l'utilisateur grâce au paradigme de l'apprentissage fédéré.

MOTS-CLÉS : vie privée, services basés sur la localisation, données de mobilité, mécanismes de protection, évaluation du risque de ré-identification, unicité de la mobilité, utilité des données.

Laboratoire (s) de recherche :Laboratoire d'InfoRmatique en Image et Systèmes d'information (LIRIS)

Directeur de thèse:

| | | |
|---|---|---|
| Ben Mokhtar Sonia | Directrice de Recherche, CNRS | Directrice de thèse |
| Bouchenak, Sara | Professeur des Universités, INSA-Lyon | Co-directrice de thèse |

Président de jury :

Composition du jury :

| | | |
|---|---|---|
| Musolesi, Mirco | Professeur des Universités, Université Collège London | Rapporteur |
| Nguyen, Benjamin | Professeur des Universités, INSA Val de Loire | Rapporteur |
| Goga, Oana | Chargée de recherche CNRS, Grenoble | Examinatrice |
| Lamarre, Philippe | Professeur des Universités, INSA Lyon | Examinateur |