



Protection against re-identification attacks in location privacy

Mohamed Maouche

► To cite this version:

Mohamed Maouche. Protection against re-identification attacks in location privacy. Networking and Internet Architecture [cs.NI]. Université de Lyon, 2019. English. NNT : 2019LYSEI089 . tel-02956245

HAL Id: tel-02956245

<https://theses.hal.science/tel-02956245>

Submitted on 2 Oct 2020

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



N°d'ordre NNT : 2019LYSEI089

THESE de DOCTORAT DE L'UNIVERSITE DE LYON

opérée au sein de
I'INSA de Lyon

Ecole Doctorale N° 512
Mathématiques et Informatique (InfoMaths)

Spécialité/ discipline de doctorat : INFORMATIQUE

Soutenue publiquement le 06/11/2019, par :
Mohamed Maouche

Protection against Re-identification Attacks in Location Privacy

Devant le jury composé de :

Fiore, Marco	Habilité à Diriger des Recherches, CNR-IEIIT	Rapporteur
Palamidessi, Catuscia	Directrice de Recherche, INRIA	Rapporteuse
Castelluccia, Claude	Directeur de Recherche, INRIA	Examineur
Petit, Jean-Marc	Professeur des Universités, INSA-Lyon	Examineur
Bouchenak, Sara	Professeur des Universités, INSA-Lyon	Directrice de thèse
Ben Mokhtar Sonia	Directrice de Recherche, CNRS	Co-directrice de thèse

Département FEDORA – INSA Lyon - Ecoles Doctorales – Quinquennal 2016-2020

SIGLE	ECOLE DOCTORALE	NOM ET COORDONNEES DU RESPONSABLE
CHIMIE	CHIMIE DE LYON http://www.edchimie-lyon.fr Sec. : Renée EL MELHEM Bât. Blaise PASCAL, 3e étage secretariat@edchimie-lyon.fr INSA : R. GOURDON	M. Stéphane DANIELE Institut de recherches sur la catalyse et l'environnement de Lyon IRCELYON-UMR 5256 Équipe CDFA 2 Avenue Albert EINSTEIN 69 626 Villeurbanne CEDEX directeur@edchimie-lyon.fr
E.E.A.	ÉLECTRONIQUE, ÉLECTROTECHNIQUE, AUTOMATIQUE http://edeea.ec-lyon.fr Sec. : M.C. HAVGOUDOUKIAN ecole-doctorale.eea@ec-lyon.fr	M. Gérard SCORLETTI École Centrale de Lyon 36 Avenue Guy DE COLLONGUE 69 134 Écully Tél : 04.72.18.60.97 Fax 04.78.43.37.17 gerard.scorletti@ec-lyon.fr
E2M2	ÉVOLUTION, ÉCOSYSTÈME, MICROBIOLOGIE, MODÉLISATION http://e2m2.universite-lyon.fr Sec. : Sylvie ROBERJOT Bât. Atrium, UCB Lyon 1 Tél : 04.72.44.83.62 INSA : H. CHARLES secretariat.e2m2@univ-lyon1.fr	M. Philippe NORMAND UMR 5557 Lab. d'Ecologie Microbienne Université Claude Bernard Lyon 1 Bâtiment Mendel 43, boulevard du 11 Novembre 1918 69 622 Villeurbanne CEDEX philippe.normand@univ-lyon1.fr
EDISS	INTERDISCIPLINAIRE SCIENCES-SANTÉ http://www.ediss-lyon.fr Sec. : Sylvie ROBERJOT Bât. Atrium, UCB Lyon 1 Tél : 04.72.44.83.62 INSA : M. LAGARDE secretariat.ediss@univ-lyon1.fr	Mme Emmanuelle CANET-SOULAS INSERM U1060, CarMeN lab, Univ. Lyon 1 Bâtiment IMBL 11 Avenue Jean CAPELLE INSA de Lyon 69 621 Villeurbanne Tél : 04.72.68.49.09 Fax : 04.72.68.49.16 emmanuelle.canet@univ-lyon1.fr
INFOMATHS	INFORMATIQUE ET MATHÉMATIQUES http://edinfomaths.universite-lyon.fr Sec. : Renée EL MELHEM Bât. Blaise PASCAL, 3e étage Tél : 04.72.43.80.46 infomaths@univ-lyon1.fr	M. Luca ZAMBONI Bât. Braconnier 43 Boulevard du 11 novembre 1918 69 622 Villeurbanne CEDEX Tél : 04.26.23.45.52 zamboni@maths.univ-lyon1.fr
Matériaux	MATÉRIAUX DE LYON http://ed34.universite-lyon.fr Sec. : Stéphanie CAUVIN Tél : 04.72.43.71.70 Bât. Direction ed.materiaux@insa-lyon.fr	M. Jean-Yves BUFFIÈRE INSA de Lyon MATEIS - Bât. Saint-Exupéry 7 Avenue Jean CAPELLE 69 621 Villeurbanne CEDEX Tél : 04.72.43.71.70 Fax : 04.72.43.85.28 jean-yves.buffiere@insa-lyon.fr
MEGA	MÉCANIQUE, ÉNERGÉTIQUE, GÉNIE CIVIL, ACOUSTIQUE http://edmega.universite-lyon.fr Sec. : Stéphanie CAUVIN Tél : 04.72.43.71.70 Bât. Direction mega@insa-lyon.fr	M. Jocelyn BONJOUR INSA de Lyon Laboratoire CETHIL Bâtiment Sadi-Carnot 9, rue de la Physique 69 621 Villeurbanne CEDEX jocelyn.bonjour@insa-lyon.fr
ScSo	ScSo* http://ed483.univ-lyon2.fr Sec. : Véronique GUICHARD INSA : J.Y. TOUSSAINT Tél : 04.78.69.72.76 veronique.cervantes@univ-lyon2.fr	M. Christian MONTES Université Lyon 2 86 Rue Pasteur 69 365 Lyon CEDEX 07 christian.montes@univ-lyon2.fr

Abstract

With the wide propagation of handheld devices, more and more mobile sensors are being used by end-users on a daily basis. Those sensors could be leveraged to gather useful mobility data for city planners, business analysts and researches. However, gathering and exploiting mobility data raises many privacy threats. Sensitive information such as one's home or workplace, hobbies, religious beliefs, political or sexual preferences can be inferred from the gathered data.

In the last decade, Location Privacy Protection Mechanisms (LPPMs) have been proposed to protect user data privacy. They alter data mobility to enforce formal guarantees (e.g., k-anonymity or differential privacy), hide sensitive information (e.g., erase points of interests) or act as countermeasures for particular attacks. In this thesis, we focus on the threat of re-identification which aims at re-linking an anonymous mobility trace to the know past mobility of its user.

First, we propose re-identification attacks (AP-Attack and ILL-Attack) that find vulnerabilities and stress current state-of-the-art LPPMs to quantify their effectiveness. We also propose a new protection mechanism HMC that uses heat maps to guide the transformation of mobility data to change the behavior of a user, in order to make her look similar to someone else rather than her past self, which preserves her from re-identification attacks. This alteration of mobility trace is constrained with the control of the utility of the data to minimize the distortion in the quality of the analysis realized on this data.

Résumé

De nos jours, avec la large propagation de différents appareils mobiles, de nombreux capteurs accompagnent des utilisateurs. Ces capteurs peuvent servir à collecter des données de mobilité qui sont utiles pour des urbanistes ou des chercheurs. Cependant, l'exploitation de ces données soulèvent de nombreuses menaces quant à la préservation de la vie privée des utilisateurs. En effet, des informations sensibles tel que le lieu domicile, le lieu de travail ou même les croyances religieuses peuvent être inférées de ces données.

Durant la dernière décennie, des mécanismes de protections appelées "Location Privacy Protection Mechanisms (LPPM)" ont été proposé. Ils imposent des garanties sur les données (e.g., k-anonymity ou differential privacy), obfusquent les informations sensibles (e.g., efface les points d'intérêt) ou sont une contremesure à des attaques particulières.

Nous portons notre attention à la ré-identification qui est un risque précis lié à la préservation de la vie privée dans les données de mobilité. Il consiste en un attaquant qui dès lors qu'il reçoit une trace de mobilité anonymisée, il cherche à retrouver l'identifiant de son propriétaire en la rattachant à un passif de traces non-anonymisées des utilisateurs du système.

Dans ce cadre, nous proposons tout d'abords des attaques de ré-identification AP-Attack et ILL-Attack servant à mettre en exergue les vulnérabilités des mécanismes de protections de l'état de l'art et de quantifier leur efficacité. Nous proposons aussi un nouveau mécanisme de protection HMC qui utilise des heat maps afin de guider la transformation du comportement d'un individu pour qu'il ne ressemble plus au

soi du passée mais à un autre utilisateur, le préservant ainsi de la ré-identification. Cet modification de la trace de mobilité est contrainte par des mesures d'utilité des données afin de minimiser la qualité de service ou les conclusions que l'on peut tirer à l'aide de ces données.

Contents

1	Introduction	3
1.1	Context	4
1.2	Privacy Threats On Mobility Data	5
1.3	Problem Statement: Re-identification	6
1.4	Countermeasures: Location Privacy Protection Mechanisms	8
1.5	Summary of Contributions	9
1.6	Thesis's Results	12
1.6.1	Publications	12
1.6.2	Developed Software	13
1.6.3	Communications	14
1.7	Structure of the Thesis	15
2	State of the Art	17
2.1	Mobility Data and Privacy Threats	18
2.1.1	Mobility Data in Diverse Form	18

2.1.2	Threats on Mobility Data	18
2.2	Re-identification Attacks	21
2.2.1	Re-identification Attacks on Mobility Data	23
2.3	Location Privacy Protection Mechanisms	24
2.3.1	LPPMs Objectives	24
2.3.2	Use Case Scenarios of LPPMs	27
2.3.3	Type of Alterations	29
2.3.4	Assessing the Effectiveness of LPPMs (Privacy/Utility Trade-off)	35
3	AP-Attack: Re-identification Attack Using Heat Maps	37
3.1	Modeling Re-Identification Attacks	39
3.2	AP-Attack Design Principles	41
3.3	Re-identification Policies	42
3.3.1	Single-Output Policy	43
3.3.2	Top-k Based Policy	43
3.3.3	Threshold Based Policy	44
3.4	Evaluation	45
3.4.1	Attack Competitors	46
3.4.2	Attacks and LPPMs Configuration	48
3.4.3	Datasets	49
3.4.4	Experimental Setup	50
3.4.5	Evaluation with Single Output Policy	50

3.4.6	LPPMs Effectiveness with Single output Policy	50
3.4.7	Evaluation with Top-K Policy	52
3.4.8	Evaluation with Threshold-based Policy	52
3.4.9	Analysis of the Parameters Affecting Re-identification	55
3.5	Conclusion	58
4	HMC: A Novel Location Privacy Protection Mechanism	61
4.1	Objectives and Roadmap	63
4.2	Recall on the Adversary Model	64
4.3	HMC Overview	65
4.4	Heat Map Alteration	66
4.5	Mobility Trace Reconstruction	68
4.6	Discussion on Alternatives for HMC	72
4.7	Experimental Evaluation of HMC	72
4.7.1	Privacy Metrics	72
4.7.2	Utility Metrics	74
4.7.3	Experimental Setup and Configurations	78
4.7.4	Privacy Evaluation	79
4.7.5	Utility Evaluation	85
4.7.6	Discussion	91
4.8	Conclusion	93

5	ILL-Attack: Mobile User Re-identification Using ERT	95
5.1	Objectives and Roadmap	96
5.2	Problem Illustration	97
5.3	Model of User Re-identification based on Small Traces	99
5.4	ILL-Attack Design Principles	101
5.4.1	Splitting Algorithm:	101
5.4.2	Data Formatting:	103
5.4.3	ILL-Attack's Classifier: The Extremely Randomized Tree (ERT)	103
5.5	Experimental Evaluation of ILL-Attack	104
5.5.1	Datasets	104
5.5.2	Experimental Setup and Configurations	105
5.5.3	Evaluation of ILL-Attack in a Session-Based Service	106
5.5.4	Evaluation of ILL-Attack in a Crowd-Sensing Application . . .	109
5.6	Conclusion	112
6	<i>Hybrid-LPPM</i>: A User-Centric Fine-Grained Multi-LPPM	113
6.1	Objectives and Roadmap	114
6.2	Problem statement	115
6.3	Design of <i>Hybrid-LPPM</i>	117
6.4	Experimental Evaluation of <i>Hybrid-LPPM</i>	118
6.4.1	Experimental Setup, Configurations and Datasets	118
6.4.2	Privacy and Utility Evaluation	119

6.5	Conclusion	121
7	Conclusion & Perspectives	123
7.1	Concluding Remarks	123
7.2	Future Work	125
	List of Figures	127
	List of Tables	131
	Bibliography	133

- Chapter 1 -

Introduction

Contents

1.1	Context	4
1.2	Privacy Threats On Mobility Data	5
1.3	Problem Statement: Re-identification	6
1.4	Countermeasures: Location Privacy Protection Mechanisms	8
1.5	Summary of Contributions	9
1.6	Thesis's Results	12
1.6.1	Publications	12
1.6.2	Developed Software	13
1.6.3	Communications	14
1.7	Structure of the Thesis	15

1.1 CONTEXT: WIDE PROPAGATION OF HANDHELD DEVICES AND LOCATION-BASED SERVICES

With the unprecedented success of handheld devices, the number of available mobile sensors is raising. According to the study of Pew Research center, 77% of the US adult population had a smartphone in 2017¹ and according to Eriksson's mobility outlook report, there are 5.1 billion mobile subscribers in 2018 and they predict growth to 7.2 billion in 2024². Many of the available applications on those smartphones ask permissions to gather the data sensed by the device and location is one of the most widely used information. As stated by Pew Research, 74% of smartphone owners used location-based services in 2013³. This industry also has an important economic impact, according to the Boston Consulting Group, the total revenue of location-based services represented \$75 billion in 2012 in the USA⁴. Hence, as stated in the "Geoprivacy Manifesto" of Keßler and McKenzie [51], location data is substantially different compared to other kinds of personal information. Since it is more and more easily obtainable due to the availability of GPS or GSM chips on handheld devices and due to the increasingly simplistic usage of available APIs to capture location data. Also, users have a substantial incentive to share their location with service providers since location improves significantly the quality of service provided by online services and enables the creation of novel services. Examples of such applications include GPS navigation and location search services such as Google Maps [37] or Bing Maps [64]). In these applications, users get directions to go to particular places and have access to a database of venues (restaurants, shops, businesses...) with their opening hours, reviews, attendance levels and so much more. Other examples include location-based social networks such as Swarm of Foursquare [25], where users can post check-ins of their movement and ask friends to join. There are location-based games such as Niantic's Pokemon GO [72] or Harry Potter Wizard Unite [73], here players are continually put into an augmented reality environment through their phone. With their movement in the real world, players can face challenges and even put their gamer-tag in particular stops, they have vanquished.

¹<https://www.pewinternet.org/fact-sheet/mobile/>

²<https://www.ericsson.com/en/mobility-report/reports/june-2019/mobile-subscriptions-outlook>

³<https://www.pewinternet.org/2013/09/12/location-based-services/>

⁴<http://www.bcg.com/documents/file109372.pdf>

In addition to interactive services, mobility data can be gathered and stored (by service providers throughout the online services or analyst throughout crowd-sensing campaigns). This data represents a great resource for city planners, businesses and researchers as it can be used for traffic information monitoring (Nericell [66]), health monitoring (PEIR [69]), social mechanisms learning (fMRI [2]) or generic research dataset gathering (APISENSE [43]).

1.2 PRIVACY THREATS ON MOBILITY DATA

However, the gathering, storage and manipulation of increasing volumes of mobility data opens several ethical and legal issues as these data are sensitive in nature and may reveal personal information about individuals. Indeed the semantic of the places the user visits could reveal sensitive information about the latter [47] such as the user's home or workplace. It could also lead to the disclosure of information the user did not wish to share. As an example, the sexual preferences of a user could be inferred from his/her regular visits to particular bars or clubs. This could lead to unwanted add targeting that could disclose the user's private life to his/her coworkers and family without the user direct consent. Other sensitive information such as one's hobbies or political alignment could also be inferred. Other threats also affect user privacy. For instance, curious adversaries could use physical interactions between users to infer social relationships between users [98]. The threat of mobility prediction also puts the user at risk [29]. For instance, the website "Please rob me" is an initiative to raise awareness about sharing our location and the disastrous consequences it can have. It uses geolocated tagged tweets to inform users about how their location sharing could inform possible robbers⁵.

With those increasing risks, users are more and more concerned about their location privacy issues. In the work of Staiano et al. [95], 60 volunteers were asked to give price to their personal data during 6 weeks (October 2013 to November 2013). Their median price for all the 6 weeks of mobility was at 22.5 euro with an opt-out percentage of 16.67%. These values are higher than all the other considered categories (communication logs, application usages and photo shot records with

⁵<http://pleaserobme.com/why>

respectively 15, 20 and 5 euros with opt-out percentages of respectively 3.34%, 0% and 8.34%). In the work of Cvrcek et al. [19], they have asked 1200 people of five European countries how much they would sell one month of mobility data (cell tower position every five minutes for one month), they received medians of approximately 50 euro for academic use and 100 euro for commercial use. We have to be careful with those surveys. Since different values can be found in different surveys. Nevertheless, they illustrate the importance that users give to their location data. In addition to the users' awareness, legislation is evolving to protect users' privacy. For instance in Europe, starting May 2018, the General Data Protection Regulation (GDPR)⁶ has been enforced. It compels every organization collecting or processing data from EU residents to be responsible and accountable for the way they manage personal data. Every organization is required to integrate privacy preserving measures by design and by default. This regulation also gives proper means for authorities to take action against the non-GDPR compliant organizations with fines up to 4% of their worldwide revenue.

1.3 PROBLEM STATEMENT: RE-IDENTIFICATION

To protect user privacy, the first possible solution considered is the full anonymization of data. Meaning that the user hides her ID contained in the data (e.g., with the usage of onion networks such as TOR). Unfortunately, this is not sufficient as mobility data acts as a quasi-identifier by itself. In the work of De Montjoye et al. [20], they studied the uniqueness of month of mobility traces of one and a half million users and they showed that with only four records, they are able to distinguish 95% of the users. This study shows that mobility data is discriminant by itself and works as a fingerprint. This means that anonymizing data after years of freely sharing location data put users at risk of being "re-identified", thus making simple anonymization ineffective at protecting users' privacy.

Re-identification attacks are the main focus of this thesis. These attacks aim at linking anonymous mobility (anonymized) to their IDs using past mobility. The first objective of this thesis is thus to study the capabilities of re-identification and to

⁶<https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX:32016R0679>

find the possible vulnerabilities on mobility data. In Equation 1.1, we define the re-identification problem considered in this thesis. The adversary has access to a set of known (i.e., with ID) traces \mathbb{K} and aims at linking an anonymous trace T (from a set of anonymous traces \mathbb{A}) to the correct known trace (thus its ID) from \mathbb{K} .

$$\begin{aligned} \mathcal{A} : \mathbb{A} &\rightarrow \mathbb{K} \\ T &\mapsto \mathcal{A}(x) = U \end{aligned} \tag{1.1}$$

In the literature, to re-identify anonymous traces, in the work of Krumm [52], they extracted users' homes. Golle and Partridge [33] showed that the pair home/work is even more discriminative. Then Gambs et al. [30] went even beyond and constructed Markov Chains between users' points of interest (i.e., particular locations the user spends a lot of time in) in order to discriminate users and to be able to assign anonymous traces to the correct users. Those attacks do not fully lighten the risk of re-identification since the characteristics these methods are built on (home, work or points of interest in general) could be hidden (e.g., by avoiding to share co-located points [83]). One might think that discriminative characteristics of mobility traces could be hidden easily but we argue that this is not the case. Hence, the first challenge of this thesis.

Challenge 1: What are the vulnerabilities of mobility data and protection mechanisms facing re-identification?

Also, those methods generally need a certain amount of mobility data. If the re-identification of anonymous traces is solely based on the recognition of the user home or workplace, does a mobility trace that does not contain these important locations always protected against re-identification? We argue that this not the case. We also argue that short traces that are not able to produce profiles such as the Mobility Markov Chain are not insensitive to re-identification.

Challenge 2: How sensitive are short mobility traces against re-identification attacks?

1.4 COUNTERMEASURES: LOCATION PRIVACY PROTECTION MECHANISMS

In order to protect users' privacy beyond simple anonymization, several protection mechanisms called Location Privacy Protection Mechanisms (LPPMs), have been proposed in the literature. The role of an LPPM is to apply data transformations to raw mobility data in order to enforce privacy guarantees to the users. These guarantees can be either well known theoretical properties, such as differential privacy [22] which bounds the knowledge an attacker may acquire when having access to the data or k -anonymity [91] which ensures that locations (or traces) are co-located in groups of at least k users. LPPMs can also use more practical techniques that either hide sensitive information (e.g., erasing points of interest with Promesse [83]) or represent countermeasures to particular attacks [63].

To reach these objectives LPPMs operate on raw mobility data at various levels of granularity. They may act at the level of individual points (e.g., Geo-I adds noise to individual geo-located coordinates [4]); they may act on a set of nearby points (e.g., Promesse removes clusters of points that correspond to user stops [83]); and they may act at the level of a sub-trace (e.g., W4M enforces k -anonymity by forcing k user traces to be co-located inside the same cylinder [1]). However, most of the existing LPPMs do not reason on the users' mobility as a whole considering multiple traces over a period of time (macroscopic vision). This limitation opens the door to powerful user re-identification attacks that try to discriminate users by reasoning on their overall mobility.

Challenge 3: How to design an effective protection mechanism that reasons on the user mobility on a macroscopic level?

Another important aspect of privacy preserving mechanisms is the privacy/utility trade-off [13]. Indeed, often when sharing obfuscated data to the service provider this results in a loss in the quality of service. For instance, if a user searches for restaurants nearby, if the location of the user is moved, the distance between her obfuscated location and her real location would alter the list of restaurants proposed. Also when sharing data in a crowd-sensing campaign, the precision of the conclusion made on the data gathered is affected by alteration that the data might go through

for privacy preservation. For instance, if the city wants to detect the most crowded places in its city but the users do not report their location in particular places, the conclusions made by the city might be altered. Thus, it is essential when designing an LPPM to evaluate its effects on the utility of the data. A protection mechanism can always scramble the data and get 100% privacy but with 0% utility. This is why privacy gains or guarantees are always relative to the utility still offered by the data.

Challenge 4: How to protect a user against re-identification attacks while maintaining the data utility?

1.5 SUMMARY OF CONTRIBUTIONS

The thesis's contributions are in two folds: **(1)** Stress current systems and find vulnerabilities. **(2)** Design countermeasures to the vulnerabilities found while taking into consideration the utility of the data. In this thesis, we first focus on designing attacks that are able to prove that mobility data is putting users at risk and that are able to show that current systems are not able to protect against re-identification attacks. From the vulnerabilities found, we designed a method based on the modification of the user behavior that protects mobility data for crowd-sensing campaigns or open data initiatives. We describe in this section the contributions of the thesis and the following chapters will go deeper into each contribution and present experiments on real mobility data.

(C1) AP-Attack: Constructing a Re-identification Attack with Robust Profiling of Users based on Heat Maps

We first start by focusing on constructing re-identification attacks. Considering an obfuscated mobility dataset and a set of user profiles learned from users past mobility, a user re-identification attack tries to re-associate a portion of the obfuscated data to its originating user (its identity). Attacks of the literature construct profiles based on points of interest that are easily hidden from adversaries. This is why in order to find vulnerabilities of systems (i.e., sensitivity to re-identification attacks), we first propose in this thesis AP-Attack (All Points Attack) a novel attack in which a user profile is represented by a heat map, a spatial aggregation of a user mobility trace in

square regions of the map.

We also propose a novel paradigm of re-identification attacks where the attacker does not consider only one possible identity as an output of a re-identification attack but rather considers multiple identities depending on a re-identification policy. The goal of a policy is to have a selection of a small number of identities that need to be further investigated. The attacker aims at having the smallest set of possible identities while including the correct identity. In consequence, we propose new ways to measure the strength of an attacker by considering the set size of possible identities and the number of false positives.

(C2) HMC: a Utility Constrained LPPM for Crowd-Sourcing and Data Publishing

Various LPPMs have been proposed in the literature. They either enforce some formal privacy guarantee (e.g., k-anonymity or differential privacy) or hide sensitive information (e.g., Promesse hides POIs) but do not explicitly protect against re-identification attacks. We propose HMC (for *Heat Map Confusion*), an LPPM that protects users against re-identification attacks by reasoning on their mobility as a whole, captured using heat maps. Specifically, in order to protect a mobility trace, HMC first uses background mobility traces of multiple users and aggregates their mobility into a single heat map per user. Then, HMC alters the mobility trace's heat map by making it look similar to the heat map of another user. To limit the decrease in data utility, HMC uses the heat map of the closest user as a basis for performing the alteration. Finally, HMC transforms back the altered heat map to a mobility traces by trying to retain as much as possible the original trace unchanged. The result is a protected mobility trace on which an attacker that runs user re-identification attacks fail in distinguishing between users.

(C3) ILL-Attack: Re-identification on Short Traces using Multi-Trace Learning rather than Profiling

We argue that evaluating the risk of re-identification when sharing data is important for the design of strong privacy preserving mechanisms. In this chapter, we propose

ILL-Attack a new re-identification attack that detects the vulnerabilities of re-identification for traces that are less sensitive to profile-based attacks. Indeed, ILL-Attack apprehends differently re-identification. it does not use mobility data to construct user profiles. Since this type of profiles demands large mobility traces to be applied. However, for shorter mobility traces (in the order of minutes or few hours), the attacker learns from multiple short behaviors in order to be able to recognize them at re-identification. ILL-Attack uses Extremely Randomized Trees to learn users' identity based on their mobility. This attack instantiates a new model of re-identification that divides the mobility traces into multiple shorter mobility traces to learn different behaviors of users in order to be able to re-identify in various scenarios. Its strength is that it can be applied to use cases of smaller length.

(C4) *Hybrid-LPPM*: A User-Centric Fine-Grained Multi-LPPM

After analyzing the results of re-identification attacks more thoroughly, we notice that users are affected differently by the LPPMs and that even portions of mobility data of the same user are not equally protected by the same LPPM. We propose to consider the particularity of each behavior of the user and design LPPMs that change their obfuscation depending on the sub-trace. This is why, we propose to make use of off-the-shelf state-of-the-art LPPMs and apply the best one for each sub-trace.

We propose *Hybrid-LPPM* that operates in a crowd-sensing application where the user goes through a privacy proxy each time it needs to send a mobility trace to the analyst. The privacy proxy is first responsible for hiding the ID of the user and hiding the source of the data. Also, it uses background knowledge sent by different users to choose the best LPPM to apply for this particular sub-trace. It also operates in a data publishing scenario, the publisher should then use *Hybrid-LPPM* instead of experimenting on each LPPM individually.

The best LPPM is chosen using two criteria: (1) Privacy: we choose the set of LPPMs that protect the most against a set of re-identification attacks trained with the gathered background knowledge. (2) Utility: from those LPPMs, we select the one with the best utility according to a chosen metric (we evaluate how much the data has been distorted).

1.6 THESIS'S RESULTS

1.6.1 Publications

International Journals

- Mohamed Maouche, Sonia Ben Mokhtar, and Sara Bouchenak. HMC: robust privacy protection of mobility data against multiple re-identification attacks. *IMWUT*, 2(3):124:1–124:25, 2018. doi: 10.1145/3264934. URL <https://hal.archives-ouvertes.fr/hal-01954041/document> [Presented at Ubicomp 2018]

International Conferences

- Mohamed Maouche, Sonia Ben Mokhtar, and Sara Bouchenak. Ap-attack: A novel user re-identification attack on mobility datasets. In *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, 2017. doi: 10.1145/3144457.3144494. URL <https://hal.archives-ouvertes.fr/hal-01785155/document>
- 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, 2018. doi: 10.1109/ICDCS.2018.00091. URL <https://hal.archives-ouvertes.fr/hal-01784557v2/document>

National Conferences and Workshops

Those conferences have peer reviews but no proceedings.

- Mohamed Maouche, Sonia Ben Mokhtar, Sara Bouchenak. Attaques de ré-identification des utilisateurs à partir de leurs traces de mobilité. *Compas*

2017

- Mohamed Maouche, Sonia Ben Mokhtar, Sara Bouchenak. HMC : Préservation de la vie privée des utilisateurs sur les données de mobilité par la protection contre les attaques de ré-identification. Compas2018
- Jugurta Ikherbane, Mohamed Maouche, Sonia Ben Mokhtar, Sara Bouchenak. Calcul multiparti et sécurisé basé sur un environnement d'exécution sécurisée. Compas 2018
- Mohamed Maouche, Sonia Ben Mokhtar, Sara Bouchenak. SFERA: Assessing Location Privacy with Re-Identification Attacks. APVP 2017
- Vincent Primault, Mohamed Maouche, Antoine Boutet, Sonia Ben Mokhtar, Sara Bouchenak, Lionel Brunie. How to Make Privacy Experimentation Open and Easy? APVP 2018
- Besma Khalfoun, Mohamed Maouche, Sonia Ben Mokhtar, Sara Bouchenak: Mood: MObility Data Privacy as Orphan Disease. Compas 2019

Ongoing Submissions

- [Accepted] Besma Khalfoun, Mohamed Maouche, Sonia Ben Mokhtar, Sara Bouchenak: Mood: MObility DataPrivacy as Orphan Disease. Middleware 2019.
- [To Submit] Mohamed Maouche, Sonia Ben Mokhtar, Sara Bouchenak: ILL-Attack: Mobile User Re-identification Using Extremely Randomized Trees.
- [Review Pending] Mohamed Maouche, Sonia Ben Mokhtar, Sara Bouchenak: User Re-identification Attacks On MobilityData: Towards a Multi-Policy Approach. IEEE TDSC.

1.6.2 Developed Software

- SFERA: A toolkit to experiment on re-identification attacks on mobility traces.
<https://github.com/mmaouche-insa/SFERA>

- HMC: A toolkit to test the Location Privacy Protection Mechanism HMC.
<https://github.com/mmaouche-insa/HMC>
- ILL-Attack: A toolkit to test ILL-Attack.
<https://github.com/mmaouche-insa/ILL-Attack>.
- Participation in Accio (main contributor is Vincent Primault): A scientific workflow management tool used to study location privacy.
<https://privamov.github.io/accio/>

1.6.3 Communications

The list of communications is listed in Table 1.1.

Table 1.1: *List of communications during the thesis*

Event	Data	Location	Title
UbiComp'18	October 10, 2018	Singapore	HMC: Robust Privacy Protection of Mobility Data Against Multiple Re-Identification Attacks
Compas'18	July 5, 2018	Toulouse, France	HMC: Privacy Protection of Mobility Data Against Multiple Re-Identification Attacks Using Macro-Mobility
IRIXYS Workshop	June 25, 2018	Gargnano, Italy	HMC: Privacy Protection of Mobility Data Against Multiple Re-Identification Attacks Using Macro-Mobility
Workshop Security Franco-Americain	December 11, 2017	Lyon, France	Protecting Users Against Re-identification Attacks Using Heat Map Alteration
IRIXYS Workshop	November 30, 2017	Hendaye, France	Protecting Users Against Re-identification Attacks Using Heat Map Alteration
MobiQuitous'17	November 10, 2017	Melbourne, Australia	A novel AP-Attack Users Re-Identification Attack on Mobility Datasets
IRIXYS Summer School	July 21, 2017	Chiemsee, Germany	SFERA: Assessing Location Privacy with Re-identification Attacks
Compas'17	June 28, 2017	Sophia Antipolis, France	SFERA: Assessing Location Privacy with Re-Identification Attacks
APVP'17	June 19, 2017	Autrans, France	SFERA: Assessing Location Privacy with Re-identification Attacks
LIRIS Security Workshop	May 30, 2017	Lyon, France	Assessing Location Privacy with Re-identification Attacks
GDR RSD ASF Winter School	March 9, 2017	Pleynet, France	Quantifying Location Privacy Using Re-identification Attacks
IRIXYS Workshop	November 1, 2016	Lyon, France	Quantifying Location Privacy Using Re-identification Attacks

1.7 STRUCTURE OF THE THESIS

The thesis is structured as follows. First, in Chapter 2, we present a state of the art on location privacy. We discuss privacy threats with more examples and we review the literature on LPPMs. Then in Chapter 3, we present AP-Attack a novel re-identification attack that uses heat maps to profile users and we present a general modeling for re-identification attacks and how to measure their effectiveness. HMC a novel LPPM designed against re-identification attacks in crowd-sensing scenarios is presented in Chapter 4. In Chapter 5, we introduce a more effective method to re-identify short traces with ILL-Attack based on learning short behaviors of users. And in Chapter 6, we present a method to make use of off-the-shelf state-of-the-art LPPMs to protect against re-identification attacks with high utility. Finally, the thesis is concluded in Chapter 7.

- Chapter 2 -

State of the Art

Contents

2.1	Mobility Data and Privacy Threats	18
2.1.1	Mobility Data in Diverse Form	18
2.1.2	Threats on Mobility Data	18
2.2	Re-identification Attacks	21
2.2.1	Re-identification Attacks on Mobility Data	23
2.3	Location Privacy Protection Mechanisms	24
2.3.1	LPPMs Objectives	24
2.3.2	Use Case Scenarios of LPPMs	27
2.3.3	Type of Alterations	29
2.3.4	Assessing the Effectiveness of LPPMs (Privacy/Utility Trade-off)	35

2.1 MOBILITY DATA AND PRIVACY THREATS

2.1.1 Mobility Data in Diverse Form

With no loss of generality, we consider all the possible locations as a metric space Λ and all the possible timestamps as a totally ordered set \mathbb{T} . A **record** is an element of $\Lambda \times \mathbb{T}$ and a **mobility trace** is a sequence of records. Thus, the set of all possible mobility traces is the free monoïde $(\Lambda \times \mathbb{T})^*$. Hence, we can make use of operations such as the concatenation of two traces and the extraction of sub-traces (similar to the operations on strings and sub-strings). Each mobility data (record or trace) is associated to one single user.

Mainly in this thesis, we consider Λ as the set of locations in the Mercator¹ projection of the earth and the timestamps as the POSIX timestamps². To simplify, we consider a record as an element of $(\mathbb{R}^2 \times \mathbb{R}_+)$. We could consider different sets of locations such as the set of cell towers or WI-FI hot-spots in a city but since the GPS representation can contain their location and all the regions they cover, we favor the Mercator representation.

2.1.2 Threats on Mobility Data

Various threats affect mobility data. In the work of Wernke et al. [100], the authors classify attacks according to the attacker's knowledge: (1) Whether the attacker has access to one record or multiple records (time constraint). (2) Whether the attacker has access to only the location data or other contextual information. They also pinpoint the target of the attacks such as the attributes in the record (*identity, location, time*). For instance, the attacker can search for the identity of the user (identity attacks), find the exact location of a user from an obfuscated location (location attacks) and the time can be used to derive information from the location such as the speed of the user or the absence of a user from her home in certain periods of time.

¹<http://mathworld.wolfram.com/MercatorProjection.html>

²https://en.wikipedia.org/wiki/Unix_time

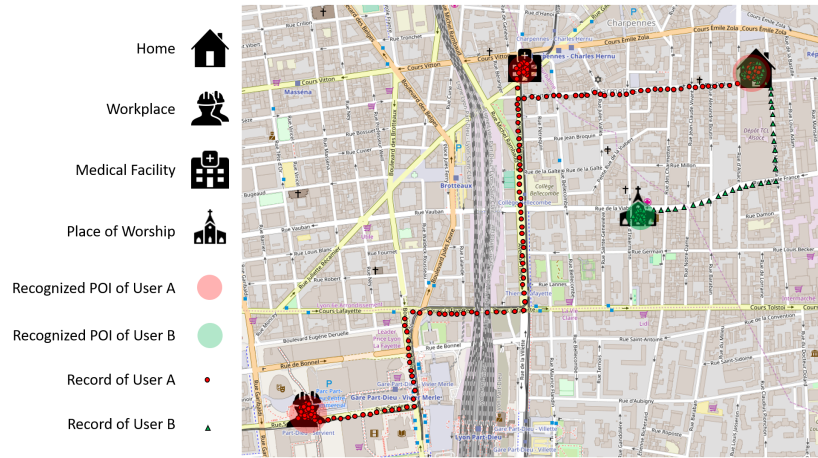


Figure 2.1: *Inference of sensitive information and social relationships between users*

Differently, in the work of Primault et al. [81], the authors categorize four threats according to the objective of the attacker. We illustrate these threats in Figure 2.1-2.3 and here is a description for each one of them.

Inference of Sensitive Information

Using mobility data an adversary can infer sensitive information about users. With the help of particular locations called *Points of Interest* (POIs) where the user spends a lot of time. Those places can be extracted easily from a mobility trace using clustering algorithms [107, 44]. Then, with those points, the attacker can use applications such as Facebook Places [23] or Google Places [36] to have the exact address, the opening hours and also the type of establishment the user visits. With this information, the attacker can infer sensitive information about the user, such as the user's home or workplace. Also, her religion if she visits worship places, sexual orientation if she visits particular bars or clubs. The attacker can infer the health status of the user if she goes through hospitals or medical facilities. In Figure 2.1, we represent the mobility of two users. From this figure, we can infer for instance user A's workplace and home.

In the work of [53], they managed to develop a system that labels places into 14 categories (e.g., home, work, transportation, place of worship, shopping, other's home...) using machine learning. Also, Huguenin et al. [47] managed to label

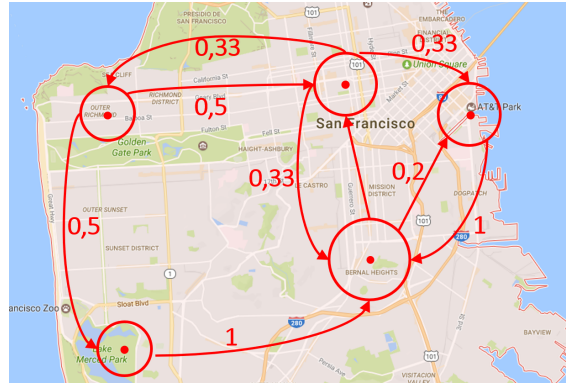


Figure 2.2: *Threat of mobility prediction Illustrated with a Mobility Markov Chain*

check-ins into 13 categories of motivation and moods of users using machine learning (e.g., "Inform about location", "Recommend it", "Appear cool/interesting", "Wish people to join me"...). In addition to the presence of a user in certain locations, the absence of records in particular times can also put the user at risk. For instance, some Redditors³ managed to extract the Muslim cab drivers from a mobility dataset [26].

Inference of Social Relationship

Using the mobility traces of multiple users, an attacker could infer relationships between users (if they visit a location at the same time). As depicted in Figure 2.1, where a home is shared by two users, a family relationship could be inferred. For instance in the work of Bilogrevic et al. [7], the authors managed to classify relationships between students (classmates or friends) using WI-FI access points. In the work of Wang et al. [98], they also infer social relationships using access-points and classify relationships into 9 categories (friends, family, neighbors, colleagues...).

Mobility Prediction

In these threats, adversaries aim at predicting future mobility of users using their past mobility. Different techniques were used, for instance, in the work of Noulas et al. [74] where they predict future check-ins using machine learning with features such as the venue's popularity, the venues' categories, the transitions between venues

³Users of a discussion website called Reddit (www.reddit.com)

and also temporal aspects. Similarly, in the work of Gambs et al. [29], they also considered the transition between locations and temporal aspects, they modeled mobility traces into Mobility Markov Chains and they used them to predict user future locations. As depicted in Figure 2.2, we have the probabilities to moving for one POI to another. In the work of Sadilek and Krumm [89], they used principal component analysis (PCA) and Fourier transformation to extract mobility pattern in order to predict mobility in a far-future (month or years).

Re-identification

Re-identification is the process of finding an identity to data that has been anonymized (identity hidden). This type of threat is the main focus of this thesis and it will be largely studied throughout the next section and the next chapters.

2.2 RE-IDENTIFICATION ATTACKS

In this section, we describe the threat of re-identification. This threat does not only exist in the context of location privacy, on the contrary, multiple types of data are at risk. We can find for instance equivalent attacks in Web Privacy [79, 97], Smart Homes Privacy [11, 24], Medical data [85, 49], Social Networks [42, 45] and even developers using their source code or binaries [99, 12]. A groundbreaking work in the theme of re-identification is the work of Narayanan and Shmatikov [71] on the Netflix Prize dataset, where they managed to match users from the anonymous Netflix dataset (containing movie ratings per user) to their public profile on another website IMDB.

The term "Re-identification" is composed of the noun "identification" which means that we distinguish an entity from others (e.g., by assigning an identifier to this entity) and the prefix "Re" which means "again", in the sense that we assign again the identifier to the entity. This means that the entity was identified but the identification was lost and we aim at assigning again its identity. The term "De-anonymization" can also be used, it means that re-identification is the counteraction of anonymizing data. Another term that is found in the literature is the term "linkability" that describes

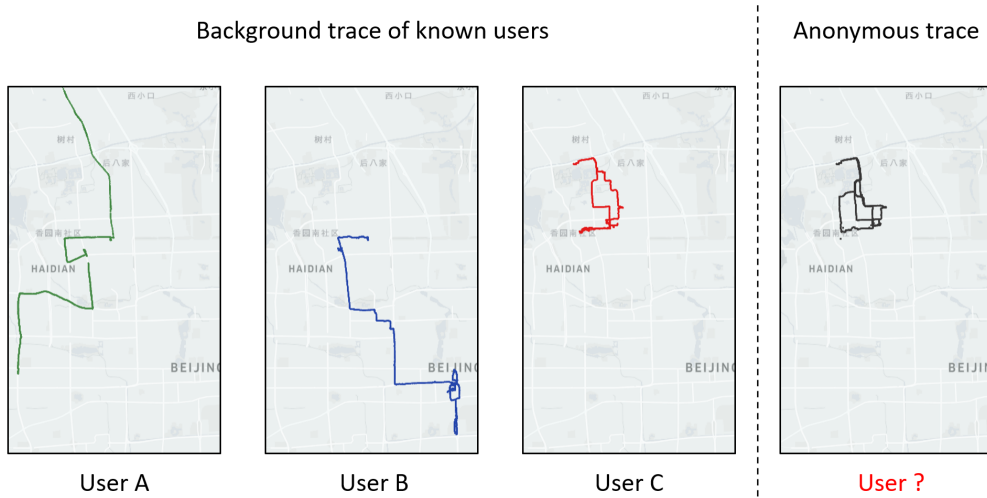


Figure 2.3: *Illustration of re-identification attacks*

the ability of data to be linked to other data. This corresponds to the capacity of an adversary to match anonymous data to a known version of it in order to identify it. In Figure 2.3, we depict the real mobility of users in Beijing (user A, B and C), we show an anonymous trace of one of the users. We notice that the anonymous trace share a lot more features with the past trace of the user C. Then, a re-identification attack would probably identify the anonymous trace as the trace of user C (it is in fact a trace of user C).

It is important also to talk about the term "Uniqueness" that is often misused in re-identification. This characterizes the property of data to be unique and thus to be "identifiable" among others. In other words, we can measure the uniqueness of a set of entities if we can find characteristics that are not shared with other entities of the dataset. While re-identification aims not only as distinguishing between entities but also to link an anonymous entity to another previous knowledge (a set of previously known identities). Studying the uniqueness of data is essential for the study of re-identification because data that is not intrinsically unique cannot be re-identified. For instance, in the famous work of De Montjoye et al. [20], they have studied months of human mobility of one and a half million individuals and they found that four GPS records are enough to uniquely identify 95% of the trace within the dataset. In uniqueness, we associate the entities of the data within the dataset itself, while in

re-identification as it is shown in this thesis, we aim at linking data from separate datasets.

Formally in Equation 2.1, we define re-identification as a function that takes as input an element of an anonymous set \mathbb{A} of entities and that finds its match from the set of known entities \mathbb{K} such as $\forall v \in \mathbb{K}$, the identity of v is known (noted $\mathcal{ID}(v)$).

$$\begin{aligned} \mathcal{A} : \mathbb{A} &\rightarrow \mathbb{K} \\ x &\mapsto \mathcal{A}(x) = v \end{aligned} \tag{2.1}$$

2.2.1 Re-identification Attacks on Mobility Data

The work of De Montjoye et al. [20] on the uniqueness of mobility data have shown that mobility of users acts as a fingerprint and thus it can be used to re-identify users. Krumm [52] looked at user de-anonymization by finding users' home addresses, they were able to find users' homes by a median error of 60 meters but the white pages system they used was not effective enough to find users' real identity. In the work of Mulder et al. [68], they use GSM data to profile users by constructing a Markov Chain between cell towers. In order to re-identify an anonymous trace, they search for the user u with the Markov Chain (i.e., the transition probability matrix $P_{x,y}^{(u)}$) that models the best the successive cell tower records of the mobility trace $(c_i)_{i=1}^n$ with the formula $\sum_{i=2}^n \log P_{c_{i-1}, c_i}^{(u)}$. In the work of Gambs et al. [30], they also use Markov Chains but between POIs (extracted from GPS mobility data). They match between anonymous traces and known users using the similarity between their respective Markov Chains. In the work of Primault et al. [82], they use POIs exclusively to re-identify users, they use a median distance between all pairs of POIs between anonymous and known traces. In this thesis, we propose a general model for this type of attacks that is presented in the next chapter (Section 3.1)

Other types of attacks that use a different paradigm than the train/test paradigm exist. Ma et al. [56] studied a type of re-identification where the anonymous traces are intercalated between the records of the known traces. Naini et al. [70] used a map grid to compare between users on a closed system and tried using a bipartite graph matching to associate traces to find identities. Some works such as Srivatsa and Hicks [94] used social network data as a side-channel to re-identify users. Specifically,

they used a contact graph identifying meetings between users extracted from a set of traces and then used a correlation with a social network graph to match users mobility with their social network account.

2.3 LOCATION PRIVACY PROTECTION MECHANISMS

In order to mitigate location privacy threats, Location Privacy Protection Mechanisms (LPPMs) have been introduced in the literature. LPPMs operate alterations on raw mobility data in order to preserve user privacy. LPPMs can be applied either record by record or on the whole mobility trace. More formally, we define a protection mechanism \mathcal{L} in Equation 2.2, it takes as input a mobility trace T and a set of parameters Υ and produces an obfuscated version of the mobility trace as an output.

$$\begin{aligned} \mathcal{L} : (\mathbb{R}^2 \times \mathbb{R}_+)^* &\rightarrow (\mathbb{R}^2 \times \mathbb{R}_+)^* \\ T &\mapsto \mathcal{L}(T, \Upsilon) = T' \end{aligned} \quad (2.2)$$

If an LPPM alters mobility data record by record, we consider its alteration of the whole mobility trace as the alteration of all the records individually. This inverse is rarely true. An LPPM that alters a sequence of records does not correspond to its application on each record individually.

2.3.1 LPPMs Objectives

LPPMs alter mobility data in order to preserve user "privacy". The notion of privacy is quite abstract and can be applied in a broad way. From the review of the literature, we notice that LPPMs aim at preserving privacy by using three approaches that are not mutually exclusive.

Enforce Formal Guarantees

These approaches aim at enforcing properties on the data to increase the privacy preservation of the user. Two main concepts are found in the literature, they are extensions of concepts from the database privacy community.

k-anonymity is a property that was introduced by Samarati and Sweeney [90] and further described in Sweeney [96], a database satisfies this property if for every subset of quasi-identifier attributes, it exists at least k entries in the database. Quasi-identifiers are special attributes that permit an attacker to link an entry on different datasets. This property can be easily generalized for the protection of location records where coordinates and time can be considered as attributes and by constructing cloaking areas (further presented in Section 2.3.3). In the work of Bettini et al. [6], they defined the property of Historical k-anonymity enforceable on mobility traces. Machanavajjhala et al. [57] proposed *l-diversity* an extension of k-anonymity, where in addition to ensuring the anonymity of the entries in a database, it ensures that sensitive attributes have at least l well represented different values for each set of quasi-identifier attributes. It was also considered in location privacy with *location diversity* of Xue et al. [104] where they enforce the semantic diversity of locations inside cloaking regions.

Differential Privacy ensures that the result of an aggregate query over a table should not be significantly affected by the presence or absence of one single element of this table [22]. More formally, A randomized mechanism \mathcal{K} that answers queries on a dataset D satisfies ϵ -differential privacy if for all datasets D_1 and D_2 differing on at most one element and for all subsets of outputs $S \subseteq \text{Range}(\mathcal{K})$ the property of Equation 2.3 is satisfied.

$$P[\mathcal{K}(D_1) \in S] \leq e^\epsilon P[\mathcal{K}(D_2) \in S] \quad (2.3)$$

As we can see from this formula, differential privacy ensures a bound on what the adversary may learn about an individual in the dataset by knowing the result of the query. This property has been extended to mobility data with Geo-Indistinguishably (Geo-I) [4]. If we consider an obfuscation technique that upon receiving a location x outputs a location y with probability k_{xy} and considering a distance d between locations, then the mechanism K associated to the probabilities k_{ab} satisfies ϵ -geo-indistinguishably if for any locations x, y, z the property of Equation 2.4 is satisfied.

$$\frac{k_{xy}}{k_{zy}} \leq e^{\epsilon d(x,z)} \quad (2.4)$$

From this formula, we can see that for an adversary x is not distinguishable from any other location z within a radius $d(x, z)$. The distinction between x and z

increases with the increase of $\epsilon d(x, z)$. In other words, the closer we get to the real location x the less information we have. For instance, an attacker may know that the user is in Lyon rather than Bruxelles and have more confidence that the user is in the Confluence district rather than being in Croix-Rousse but the adversary cannot know the exact location of the user. Thus making this property strong against Location attacks. In practice, adding a two dimensional Laplacian noise to the data is sufficient to enforce Geo-Indistinguishability on a record.

Hiding Sensitive Information

Some LPPMs alter data to hide sensitive information about the user. For instance, *Promesse* of Primault et al. [83] erases POIs with a time-distortion and a speed smoothing technique. In other LPPMs, they avoid sharing particular sensitive information such as in *SRide* [3] where origin and destination of ride-shares are encrypted to hide them from the service provider and the compatibility between clients and drivers is computed with the help of homomorphic encryption and secure multi-party computation. In the work of Xu and Cai [103], the user specifies public regions where she would feel comfortable to be reported in. In the work of Riboni and Bettini [86] before publishing a dataset of check-ins, they first filter check-ins that are in regions where the user did too many check-ins since it may be a sensitive region for the user.

Counter-Attack

Some LPPMs protect users against particular attacks or particular categories of attacks. For instance, with Geo-Indistinguishability the user bounds the quality of location attacks, when enforcing k -anonymity it bounds the probability of linking one correct to its user to a probability of k^{-1} . If you erase POIs you ensure that an attacker cannot find easily the place visited by the user and hence you limit the personal sensitive knowledge an attacker can have. Cryptography can also be a strong tool in these types of protection. For instance, in the work of Mascetti et al. [63], they propose to use cryptography to find nearby friends without disclosing the users' locations.

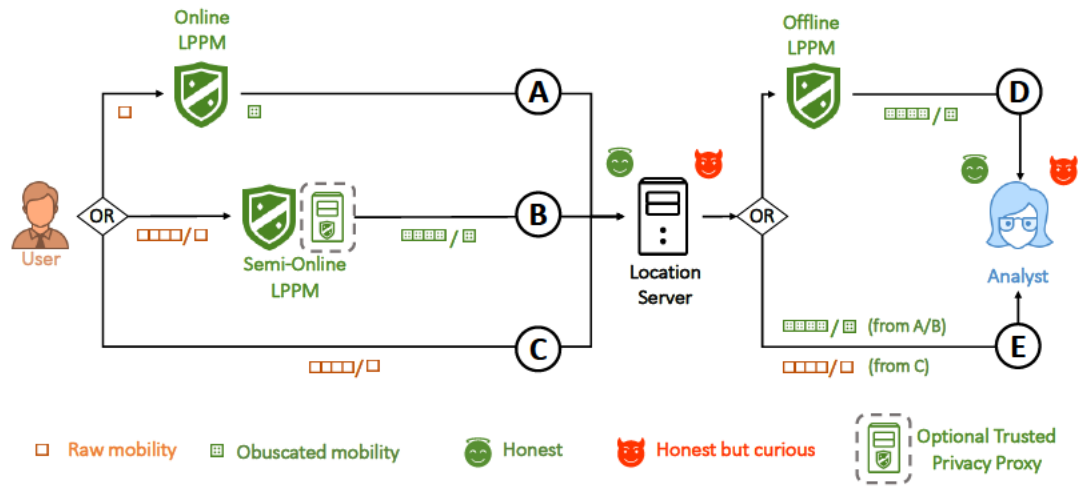


Figure 2.4: *Different use case scenarios for the usage of LPPMs*

2.3.2 Use Case Scenarios of LPPMs

In Figure 2.4, we depict the various scenarios in which a user can interact with an LPPM. The two main ones are first the **online scenario** where the data is transformed record by record on the user device (path A in the figure) and the second one is the **offline scenario** where the data publisher transforms all the mobility traces of all the users before publishing the dataset (path D). We also consider an intermediate scenario where the user can go through a privacy proxy in order to protect her data (path B). We call it **semi-online** since it does not respect the real-time paradigm of the full online scenario. It either delays records to apply transformations on a sequence of records or it interacts with a privacy proxy that can have access to the records of other users and act accordingly. These scenarios are not mutually exclusive since a user could protect her mobility data with an LPPM on her local device (online or semi-online) and interact with a location-based service that would later publish the dataset after the application of another LPPM (offline).

In Table 2.1, we compare the three scenarios on different characteristics. The main difference between the three scenarios is the constraint we put on the LPPMs. In the online scenario, an LPPM applies its alteration without the knowledge of

Table 2.1: *A Comparative between different use case scenarios of LPPMs*

Characteristics	Scenario		
	Online	Semi-online	Offline
Usage	Location-based Service	Crowd-Sensing and Data Publishing	Crowd-Sensing and Data Publishing
Applicable on Individual Records (Real-time)	By Definition	Possible	Possible but not useful
Applicable on a Sequence of Records	Incompatible	Possible	Possible
Trust	Trust the Service Provider	Can rely on a Privacy Proxy	Trust the Data Publisher
Cohesive Users	Through P2P	Through the Privacy Proxy	Complete Knowledge of All Users

future records. In the offline scenario there are no constraints on the LPPM, it can alter a whole mobility trace and even consider the interaction between users. This is why, we introduce the semi-online scenario, an intermediate scenario where the LPPM can consider a batch of records and may delay its alteration to receive data of other users without knowing fully all data that would be received further from all the users (including the user currently being protected by the LPPM). Hence, it also modifies the usage of those LPPMs, an online scenario is adapted to a use case where the user needs to protect her individual records one by one before sending it to a location-based service, while, in use cases where multiple records need to be protected as a group, the user might need more of an offline or semi-online scenario. The LPPMs also differ on whether or how the users interact with each other. In an offline scenario, we can consider a full user interaction while for an online scenario if the user needs to interact outside of the location server surveillance they need to either use a P2P network linking them or trust a third party such as a privacy proxy (semi-online).

With regards to the design of LPPMs, it is important to notice that the constraints are incremental with the order: offline (fewer constraints) \rightarrow semi-online \rightarrow online. While with regards to the usage the order is inverted, every LPPM that is applicable the online scenario can be applied in the semi-online and offline scenario and every



Figure 2.5: *Illustration of various types of mobility trace and record alteration*

LPPM that is applied in the semi-online can be applied in the offline scenario. To illustrate with some examples. In the work of Huang et al. [46] each record is individually transformed into three dummy records, which do not depend on other users or on the position of future records produced by the user. This LPPM can then be applied in an online scenario directly on the user device with interactivity. On the other hand an LPPM such as Promesse [83] that applies a time distortion mechanism on a batch of records cannot be applied in an online way but rather in a semi-online fashion. As for W4M [1], since in addition to modifying a whole mobility trace it needs to have access to all the data of the other users, this LPPM needs to be applied in an offline scenario.

2.3.3 Type of Alterations

In order to mitigate the location privacy threats, LPPMs apply various transformations to the mobility data. In this section, we categorize the different types of alterations found in the literature. We consider both LPPMs that alter mobility data record by record or in the form of a batch of records. As an illustration, we draw examples of alterations for each type of alteration in Figure 2.5.

Generalization-based

In this alteration methods, instead of sending the exact location of a record, the LPPMs send a region surrounding the location. Which means that the location is less precise but still correct. This region can either be represented by a well-established sub-division of the map (e.g., street, neighborhood, city, area code ...) or by some kind of ad-hoc region generated by the LPPM. The LPPM can either send the boundaries of the zone or one coordinate that represents the zone (the center of the zone for example). In Figure 2.5a, we show the records of three users being obfuscated, the LPPMs would send for the users the location R_c instead of the GPS coordinates. These methods have been successfully used to enforce k-anonymity by creating cloaking areas (introduced by Gruteser and Grunwald [40]) in which at least k users are present. For instance, Gedik and Liu [31] propose *Clique Cloak* a method where a trusted proxy coordinates different users by delaying queries in order to enforce k-anonymity and sends cloaking areas. Mokbel et al. [67] proposed *Casper* an enhancement of this method that is scalable to arbitrary levels of k and that uses randomization in order to prevent the reverse engineering of the cloaking areas. The authors even proposed a P2P version [18]. An interesting work is a "feeling-based" method of Xu and Cai [103], the user specifies public regions where she would feel comfortable to be reported, which is more expressive for users since they do not need to apprehend formal privacy concepts such as k-anonymity. Most of the methods presented until now are online or semi-online methods, they mostly reason on records alone. Some work considers full mobility traces such as Gramaglia and Fiore [38] with *Glove*, they use a similarity metric between trajectories to merge traces with generalization zones to provide k-anonymity. Gramaglia et al. [39] also proposed $k^{\tau, \epsilon}$ -anonymity an extension of k-anonymity that includes temporal information. This privacy constraint considers an attacker that may have access to a previous mobility of the users during a period of time of at most τ and it allows the attacker to eavesdrop on the user for a period of time of at most ϵ .

These methods allow users to hide their precise location but unfortunately it is not enough to hide the discriminative patterns of users. A user might always be surrounded by other users but few users would actually follow her around the multiple regions she visits and the combination of different regions visited discriminates the user and makes her vulnerable to re-identification. Also, if those methods aim at

enforcing k-anonymity, they need to be applied in systems and applications where a huge amount of users is active. It cannot be applied for crowd-sensing campaigns where users might find themselves alone. Since the cloaking areas, in this case, would be huge and thus providing no use for the analyst.

Perturbation-based

In this alteration methods, the mobility data is perturbed in the sense that for a given record $r = (x, y, t)$ the record is either erased or moved. The perturbation can be either in space or in time (i.e., $r' = (x + \Delta x, y + \Delta y, t + \Delta t)$). As depicted in Figure 2.5b, we show a record being obfuscated by a perturbation method. This type of alteration is useful to protect users against location attacks that aim at finding the user exact location. The perturbed record is still useful for multiple applications where it is not necessary to have access to the precise location of the user but a close record is enough (for instance finding surrounding POIs). When the noise added to a record is calibrated it can become the basis of methods that enforces some formal guarantees such as differential privacy [22].

For instance, Andrés et al. [4] proposed an adaptation of differential privacy to location data with *Geo-Indistinguishably* (Geo-I for short). By adding 2-dimensional Laplacian noise to a mobility record x it ensures that the produced record y is ϵ -Geo-indistinguishable. This means that for every possible location z , the probability of x being the real location is bounded by the distance between z and x . As summarized in the overview of Chatzikokolakis et al. [17], if a protection mechanism has a probability k_{xy} of producing y from a real location x then for any location z this following property is verified $k_{xy} \leq k_{zy}e^{\epsilon d(x,z)}$. They also proposed an extension called "elastic distinguishability" [16] that manages the noise differently depending on the area where the users have been obfuscated, they consider that a location is better protected in an area with a high "privacy mass", which means that the area is rich in venues, for their experiment they used the POI data of Open Street Map data [62] to evaluate this privacy mass.

An important issue in the management of the noise is the management of the ϵ budget, Geo-I as differential privacy is composable, meaning that applying n times ϵ -Geo-I on records will result in a $n\epsilon$ -Geo-I protection (higher the ϵ worse is the privacy).

In order to mitigate this issue, Chatzikokolakis et al. [15] proposed a solution that does not systematically add noise to each location. It rather simulates the results of a location attack based on the previously reported location and depending on its success outputs either the prediction in case of success or the real location with noise which decreases the privacy budget. The solution presented previously apply transformations to the data without considering previously reported location, Xiao and Xiong [101] [102] proposed a mechanism that considers a convex hull of the most probable locations at each timestamp and they use a Hidden Markov Chain to model time correlation to enforce a temporal aware differential privacy property.

The methods presented above are mainly applied in online or semi-online scenarios. Other methods have also been proposed to publish datasets in an offline scenario. For instance, Mir et al. [65] proposed a method to generate synthetic CDRs (call detail records) with differential privacy called *DP-WHERE* (a differentially private extension of *WHERE* [48]). They start by building a model on real CDRs data by computing several histograms and then they add noise to each of them to achieve differential privacy. The synthetic CDR can be generated by using the obfuscated histograms.

Perturbation methods that are applied on records independently are intrinsically useful against location attacks. But if the alteration does not consider the correlation between records and the patterns of the mobility then they can be ineffective against re-identification attacks as this will be shown in the following chapters.

Dummy-based

In this category, an LPPM adds fake data alongside the correct one or in place of the correct one to confuse the attacker. The added dummies can either have the same IDs as the user that the LPPM protects or can be a new dummy user created by the LPPM. In Figure 2.5c, we show a user that reports three different locations to a navigation service without sending the correct one. The first works that used dummy-based methods were simplistic by generating endpoints and producing fake mobility trace between them with random speeds or rotating real trajectories such as in You et al. [105]. In the work of Kato et al. [50], they generate fake mobility traces that considers users' stops and adds intersections between users to increase

the confusion. Shankar et al. [93] proposed *SybilQuery* an LPPM that generates fake trips suited for navigation applications. It creates fake trips that start and end in different locations but those fake trips maintain properties of the real trip such as the length and the semantics of the area of the endpoint. Bindschaedler and Shokri [8] proposed a synthetic trace generator that uses sample traces of a real mobility dataset as a seed for the generation of the synthetic dataset. This dataset is supposed to resemble human mobility and have statistical features similar to the one of the real traces without leaking significant information about any particular individual whose data is used in the synthesis process. In Huang et al. [46], each record is transformed into three random records in a region surrounding the correct one, it also ensures correct responses from location searching services through trilateration.

An important aspect of dummy-based solutions is their capacity to produce realistic data. For instance, in the work of Peddinti and Saxena [77], they managed to construct an attack that is able to find the fake queries. With a value of $k = 5$ (i.e., 4 fake queries per real query) and an attacker that has access to previous mobility data, they built a machine learning model that is able to find 93.67% of the real trips (true positive rate) with a false positive rate of 2.02%. If the attacker does not have access to past mobility, they search for correlations between trips at different iterations to erase improbable trips (using maximum speed limitations) and match between different successive trips by matching previous destinations with new sources by selecting the one with the closest speed compared to the previous trip average speed. The results show that they manage to obtain a true positive rate of 40% (twice bigger than the random 20% with $k = 5$). Another important issue with the dummy-based methods is that they increase the quantity of data to process. *SybilQuery* multiplies the number of requests sent to the LBS by its parameter k , in the work of Huang et al. [46] each record is replaced by three records which increases the load on the servers. This type of alteration is also inapplicable in some use cases, for instance, if an analysis is based on counting the presence of users in one place the addition of fake users makes the results incorrect.

Mix-zones

An area called mix-zone is designed by the LPPM either statically or dynamically. In this mix-zone, no user sends data. But when leaving the mix-zone the IDs of the

users are switched. In Figure 2.5c, we depict three users going through a mix-zone and their IDs being switched when leaving the mix-zone. It was first introduced by Beresford and Stajano [5] taking inspiration from the concept of mix-networks in routing. In addition to the main mechanisms of mix-zones, some authors studied their optimal placement [27, 55]. Different information can be used to construct the mix-zones such as POIs for Liu et al. [55], road network and speed for Palanisamy and Liu [76] or even social networks for Gong et al. [34].

Unfortunately, those methods are mostly applicable in online interactive use cases. Since, if the traces are gathered to publish a dataset, one real single trace would have multiple IDs and would be considered as coming from multiple different users, as a consequence, the use of user-centric analysis on this dataset would be compromised. Another issue is the number of mix-zones placed in the city. If they are too few some attacker might still attack the segments where the user does not change the ID and the trace could also be reassembled using re-identification attacks. On the contrary, too many mix-zones would result in too few data gathered since data is not shared when a user is inside a mix-zone.

Protocol-based

This category is not a type of "alteration" since records are not altered but rather those LPPMs are protocols specific to an application that are designed to protect the users' privacy. Most of the solution of this category make use of either cryptographic tools or secure multi-party computation principles. For instance, Mascetti et al. [63] proposed two protocols that find nearby friends by sharing cryptographic keys. *KOI* was proposed by Guha et al. [41], it is a protocol that makes use of two non-colluding servers that share parts of a request through a cryptographic scheme. The goal is to isolate the three components of a location query (namely the user id, the location trigger and the message of the query). This changes the paradigm of how applications should be built moving from a location-response paradigm to an event-centric paradigm with triggers. Aïvodji et al. [3] proposed *SRide*, a privacy preserving protocol for ride-sharing systems, they intend to hide sensitive information from the ride-sharing provider such as the origin and the destination of the rides. For this purpose, they use homomorphic encryption to find compatible drivers and they use secure multi-party comparisons to assign scores to the different compatible

drivers.

Those types of mechanisms are generally the ones that offer the best privacy/utility trade-off for the intended application but are unusable on other applications. They also often do not work on legacy systems, they ask for a renewal of the service provider's system.

2.3.4 Assessing the Effectiveness of LPPMs (Privacy/Utility Trade-off)

LPPMs intrinsically aim at optimizing a trade-off between the utility of the data and the privacy of the users. The utility can be considered in diverse form, it can be individual utility where the user evaluated the quality of service she will receive or it can be evaluated as the quality of the conclusions we can extract from all the analysis of the data. Two ways can be used: (1) service-centric: by evaluating the distortion in the service (or analysis) before and after obfuscation. (2) data-centric: by evaluating the distortion directly on the data and see the difference before and after obfuscation. The latter is more generic since it can be used to apprehend the results of all the services that need precise data according to the evaluated information. For instance, if we evaluate a spatial distortion on the data, it is a good indicator for all services that need precise locations but it does not represent a good indicator for services where time is important. The service-centric metrics are more specific and thus more precise for a particular service.

Examples of service-centered utility metrics include the measure of the quality of venue recommendation in the work of Riboni and Bettini [86]. In the work of Bindschaedler and Shokri [8], they study how a recommender system profile would be polluted by the fake queries sent to the service provider. To do so, they compute the number of distinct semantic classes of surrounding POIs of the user before and after obfuscation. Examples of utility metrics that compute the distortion in the data include the computation of spatial distortion in the work of Primault et al. [83], the authors consider the orthogonal distance between the obfuscated record and the trajectory of the original trace. In the work of Chatzikokolakis et al. [16], the authors evaluate the average expected geographical distance between the record and its obfuscated version. In the work of Gramaglia et al. [39] since they use generalization

of records they compute the size of the generalization area as a utility measure (time and spatial). The differences between the areas covered by the mobility traces before and after obfuscation are measured in the work of Cerf et al. [13].

Considering the privacy evaluation. It corresponds to the evaluation of the objectives described in Section 2.3.1. (1) If the LPPM enforces a formal guarantee, its parameter would evaluate the LPPM privacy effectiveness (e.g., k of k -anonymity, ϵ of differential privacy). (2) By evaluating what sensitive information the attacker may gather before and after obfuscation. (3) By evaluating the effectiveness of attacks before and after obfuscation. For instance, in the work of Cerf et al. [13], they search for POIs retrieval of an attack from the mobility trace before and after obfuscation. In the work of Primault et al. [83], they evaluate the effectiveness of a re-identification attack before and after obfuscation. In the remaining of the thesis, we define the metrics for both privacy and utility used in every evaluation in its corresponding section.

- Chapter 3 -

AP-Attack: Re-identification Attack Using Heat Map Profiles of Users

Contents

3.1	Modeling Re-Identification Attacks	39
3.2	AP-Attack Design Principles	41
3.3	Re-identification Policies	42
3.3.1	Single-Output Policy	43
3.3.2	Top-k Based Policy	43
3.3.3	Threshold Based Policy	44
3.4	Evaluation	45
3.4.1	Attack Competitors	46
3.4.2	Attacks and LPPMs Configuration	48
3.4.3	Datasets	49
3.4.4	Experimental Setup	50
3.4.5	Evaluation with Single Output Policy	50
3.4.6	LPPMs Effectiveness with Single output Policy	50
3.4.7	Evaluation with Top-K Policy	52
3.4.8	Evaluation with Threshold-based Policy	52

3.4.9	Analysis of the Parameters Affecting Re-identification . .	55
3.5	Conclusion	58

OBJECTIVES AND ROADMAP

In Chapter 1, we presented the outgrowing risk on users' privacy. More particularly, we focused on location data being increasingly gathered by service providers and the threats that this phenomenon opens. We established in Chapter 2 a state of the art of the threats affecting location privacy with a focus on user re-identification attacks, which are the main subject of this thesis. We also established a state of the art on the protection mechanisms (LPPMs) currently available in various forms. We argue that the threat of user re-identification attack is significant and that we should assess the effectiveness of LPPMs against these attacks.

In this chapter, we focus on measuring the user re-identification threat. Considering an obfuscated mobility dataset and a set of user profiles learned from users past mobility, a user re-identification attack tries to re-associate a portion of the obfuscated data to its originating user (its identity).

We propose AP-Attack (All Points Attack) a novel attack in which a user profile is represented by a heat map, a spatial aggregation of a user mobility trace in square regions of the map. We also propose a novel paradigm of re-identification attack where the attacker does not consider only one possible identity as an output of a re-identification attack but rather considers multiple identities depending on a re-identification policy. The goal of a policy is to have a selection of a small number of identities that need to be further investigated. More precisely, the attacker aims at having the smallest set of possible identities while including the correct identity. In consequence, we propose new ways to measure the strength of an attacker by considering the set size of possible identities and the number of false positives identities. In this chapter, we also analyze the effect of the number of users in the system on the protection mechanisms. We also evaluate the proportion of mobility an attacker needs in order to properly re-identify users.

The experiments are conducted on four real mobility datasets using three state-of-the-art LPPMs (i.e., Geo-I [4], Promesse [83] and W4M [1]). The results show that AP-Attack the attack based on heat map profiling outperforms POI-Attack and PIT-Attack. The median anonymity set size between LPPMs can vary from $k = 1$ to $k = 78$ depending on the dataset. The comparison between the different

policies show that a non-systematic method like the threshold-based policy - which selects a variable number of identities - is able to outperform the theoretical bound of a systematic method such as the Top-k policy - which systematically selects k identities - in terms of average precision and average false positive rate.

The work presented in this chapter has been published in MobiQuitous 2017 [59] and a journal extension is currently in review in IEEE TDSC journal.

Roadmap This chapter is structured as follows. First, we present in Section 3.1, a model for re-identification attacks. We present the re-identification attack AP-Attack in Section 3.2. In Section 3.3, we present different user re-identification policies. Further in Section 3.4, we evaluate AP-Attack and two state-of-the-art attacks of the literature POI-Attack and PIT-Attack against state-of-the-art LPPMs using four real datasets. We use the different re-identification policies to assess the effectiveness of the LPPMs. And we investigate two parameters that might affect the re-identification (the number of user in the system and the proportion of mobility trace eavesdropped). We conclude this chapter in Section 3.5.

3.1 MODELING RE-IDENTIFICATION ATTACKS

Let $\mathbb{U} = \{U_1, U_2, \dots, U_N\}$ be the set of users in the system. The first phase of a re-identification attack is the training phase in which the adversary builds a knowledge base about the users in the system. In real systems, this phase may correspond to a period of time where users were using a geo-located service without protecting their mobility data. This phase is depicted in the left part of Figure 3.1. Specifically, we assume that for each user U_i , the adversary has access to a set of mobility traces corresponding to her past mobility, i.e., KD_i (which stands for Known user Data). And the set of all mobility traces known by the adversary is noted $\mathbb{KD} = \{KD_1, KD_2, \dots, KD_n\}$. From each of these traces KD_i , we assume that the adversary builds a user profile $\mathcal{P}(KD_i)$ that characterizes the user mobility as depicted in the left part of Figure 3.1. This profile is specific to each re-identification attack as further discussed in Section 3.4.1.

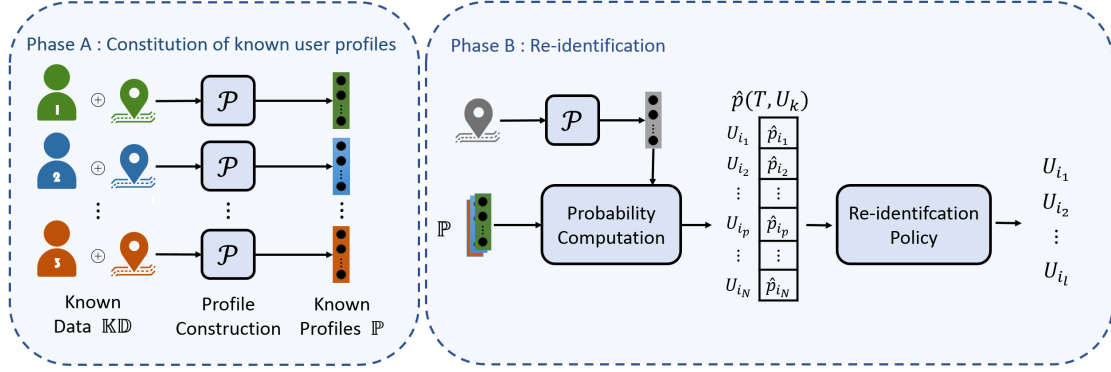


Figure 3.1: *Re-identification attacks process from collecting phase to re-identification phase*

The second phase of a re-identification attack is depicted in the right part of Figure 3.1. In this phase, we assume that the adversary obtains an anonymous trace T and then builds a profile $\mathcal{P}(T)$ with a similar structure as the one of the set of known profiles \mathbb{P} . Then, the attacker computes a similarity measure between the profile $\mathcal{P}(T)$ and each profile of $P_i \in \mathbb{P}$, i.e., $s(\mathcal{P}(T), P_i)$. Using this similarity measure, the attacker estimates (see Equation 3.1) the probability that the anonymous trace originates from a user U_i (i.e., $P[\mathcal{ID}(T) = U_i | \mathbb{KD}]$).

$$\begin{aligned} \hat{p}(T, U_i) &= \hat{P}[\mathcal{ID}(T) = U_i | \mathbb{KD}] \\ &= \frac{s(\mathcal{P}(T), P_i)}{\sum_{P_k \in \mathbb{P}} s(\mathcal{P}(T), P_k)} \end{aligned} \quad (3.1)$$

Finally, as defined in Equation 3.2, the re-identification attack \mathcal{A} outputs the list of candidates depending on its decision policy (Section 3.3)

$$\begin{aligned} \mathcal{A} : \mathbb{UD} &\rightarrow \mathbb{U}^l \\ T &\mapsto \mathcal{A}(T, \mathbb{KD}) = (U_{i_1}, \dots, U_{i_l}) \end{aligned} \quad (3.2)$$

In addition to the way user profiles are modeled, another key element for the success of a re-identification attack is the similarity metric used to compare anonymous data with known user profiles.




Figure 3.2: *From mobility trace to heat map*

3.2 AP-ATTACK DESIGN PRINCIPLES

We present in this section **AP-Attack** (All Points Attack) a novel re-identification attack that uses the whole user mobility data to form user profiles. Specifically, instead of focusing on a sub-set of points (e.g., those constituting POIs), AP-Attack aggregates all the points enclosed in a user mobility trace into a heat map structure. More precisely, as shown in Figure 3.2, the map is subdivided into a grid with cells of the same size. Then, in each cell, the number of records found in it is computed. As such, each cell will reflect the intensity of user movement in the corresponding geographical zone. This allows distinguishing between extremely, moderately, slightly frequented cells and unfrequented cells for each user. Thereby, $\mathcal{P}_{AP}(KD_i)$ returns a probability distribution where each value $\mathcal{P}_{AP}(KD_i)^{(k)}$ represents the probability that the owner of the trace U_i goes through the cell k . In order to be able to take into consideration the whole world map, the representation of each heat map is a mapping between unique cells that the user passed by and their probability. This way, each user would have a different sized map adapted to how wide her mobility was. This can also be seen as a sparse matrix.

Furthermore, we translate the distance between two profiles with the distance between two probability distributions. To compute this distance we can rely on classical distance metrics between probability distributions such as the ones surveyed in [14]. With respect to the experiments we did, one of the best metric to choose from is the Topsoe divergence defined in Equation 3.3. Where X and Y represent the list of cells in the two heat maps we compare. So X_i is the probability of the user represented by the heat map X going through the i th cell. This divergence is based on Shanon's concept of probabilistic uncertainty or entropy. It is a derived symmetric

version of the Kullback Leibler divergence [14] which measures the information deviation. This is adapted to our case since we measure how much a heat map can be used to characterize the mobility of a user that is represented by another heat map.

$$d(X, Y) = \sum_i \left[X_i \ln \left(\frac{2X_i}{X_i + Y_i} \right) + Y_i \ln \left(\frac{2Y_i}{X_i + Y_i} \right) \right] \quad (3.3)$$

Since the model presented in section 3.1 uses similarities rather than distances, we normalize the result of the distance between 0 and 1 then we compute $s \equiv 1 - d$. This attack does not take into consideration the POIs only but also mobility as a whole. Thus, it makes the LPPMs that are based on erasing POIs less effective.

3.3 RE-IDENTIFICATION POLICIES

In the model presented in the previous section, the attack outputs a list of ordered identities associated with their estimated probabilities of being the correct identity of the anonymous trace. In this section, we present three policies the attacker can choose from, to decide which identities to take into consideration. These policies are: a simple single output policy, a top-k based policy, a threshold-based policy. For each policy, we propose measures to quantify the strength of the attack and the quality of its results.

Table 3.1: *Example of attack policies*

User	Pr	Single output	Top-k $k = 3$	Threshold based $\alpha = 0.25$
U_{i_1}	0.5	U_{i_1}	U_{i_1}	U_{i_1}
U_{i_2}	0.4			
U_{i_3}	0.07			
U_{i_4}	0.02			
\vdots	\vdots			

3.3.1 Single-Output Policy

In this policy, the attacker outputs one identity that is the most probable one (Equation 3.4).

$$\begin{aligned}\mathcal{A}(T, \mathbb{KD}) &= U_a \in \mathbb{U} \\ &= \arg \max_{U_i \in \mathbb{U}} \hat{p}(T, U_i)\end{aligned}\tag{3.4}$$

Using this policy, we can measure the accuracy of an attack using the re-identification rate. For this, we use a set of anonymous traces \mathbb{UD} , the attack is re-iterated on each element UD_i . The success of an attack is then computed based on the number of correct re-associations the attack performs between anonymous traces and known user profiles. To do this, we employ an oracle \mathcal{ID} that is able to disclose for each anonymous trace UD_i its owner identity $\mathcal{ID}(UD_i) = u(UD_i)$. This way, we can compute the **user re-identification rate as follows** (Equation 3.5):

$$r(\mathcal{A}, \mathbb{KD}, \mathbb{UD}) = \frac{\sum_{UD_i \in \mathbb{UD}} \begin{cases} 1 & \text{If } \mathcal{A}(UD_i, \mathbb{KD}) = \mathcal{ID}(UD_i) \\ 0 & \text{Else} \end{cases}}{|\mathbb{UD}|}\tag{3.5}$$

In Table 3.1, we present an example of a re-identification attack. For this policy, the attacker outputs one identity U_{i_1} .

3.3.2 Top-k Based Policy

In this policy, we consider the sorted set of most probable identities (Equation 3.6). the attacker always selects the k identities with the highest probabilities (Equation 3.7).

$$\begin{aligned}\mathbb{S}(T, \mathbb{KD}) &= (U_{i_1}, \dots, U_{i_N}) \\ \forall p < q \leq N : \hat{p}(T, U_{i_p}) &\geq \hat{p}(T, U_{i_q})\end{aligned}\tag{3.6}$$

$$\begin{aligned}\mathcal{A}^{(k)}(T, \mathbb{KD}) &= (U_{i_1}, \dots, U_{i_k}) \in \mathbb{U}^k \\ \mathcal{A}^{(k)}(T, \mathbb{KD}) &\subseteq \mathbb{S}(T, \mathbb{KD})\end{aligned}\tag{3.7}$$

Using this policy, a way to measure the effectiveness of the attack is to identify the minimum k for which the real identity of the anonymous trace is part of the

output set $\mathcal{A}^{(k)}(T, \mathbb{KD})$ ($\min K$ defined in Equation 3.8). To evaluate the attack using multiple anonymous traces, we either average each $\min K$ or study the distribution of all the k (the higher the better is the privacy).

$$\min K(\mathcal{A}, T, \mathbb{KD}) = \min\{k \in \mathbb{N} \mid \mathcal{ID}(T) \in \mathcal{A}^{(k)}(T, \mathbb{KD})\} \quad (3.8)$$

In Table 3.1, we present an example of a re-identification attack for this policy with $k = 3$. The attacker outputs three identities $\{U_{i_1}, U_{i_2}, U_{i_3}\}$. Let us say that the correct identity is U_{i_2} then $\min k = 2$.

3.3.3 Threshold Based Policy

The disadvantage of the top-k policy is the risk of false positives. This is why we propose a policy that chooses users only according to a certain level of confidence. In this policy, we consider the sorted set of most probable identities (Equation 3.6). the attacker selects all the identities that have a probability above the threshold limit (Equation 3.9).

$$\begin{aligned} \mathcal{A}(T, \mathbb{KD}, \alpha) &= (U_{i_1}, \dots, U_{i_{l_\alpha}}) \\ l_\alpha &= \min \{0 \leq l \leq N \mid \forall p \leq l : \hat{p}(T, U_{i_p}) > \alpha\} \end{aligned} \quad (3.9)$$

The advantage of this policy is its ability to resize the considered group of identities. Hence, avoiding to systematically add false positives to the group. To measure the relevance of this policy, we use two metrics: (1) The first one measures the average precision of all the considered groups (Equation 3.10¹) This metric rewards finding the correct user but penalizes a policy with a too-large set size. (2) The second metric measures the rate of false positives (Equation 3.11). If the policy finds exclusively wrong users then its rate will be equal to 1, if it finds the user among l_α other users then the rate will be equal to $1 - \frac{1}{l_\alpha}$. But if the policy considers this set too uncertain and outputs no identity then this metric rewards the attacker by considering a false positive rate equal to 0.

¹We do not put α in the parameters of the metric since it may be used for other policies.

$$p(\mathcal{A}, \mathbb{KD}, \mathbb{UD}) = \frac{1}{|\mathbb{UD}|} \sum_{UD_i} \frac{\begin{cases} 1 & \text{If } \mathcal{ID}(UD_i) \in \mathcal{A}(UD_i, \mathbb{KD}) \\ 0 & \text{Else} \end{cases}}{|\mathcal{A}(UD_i, \mathbb{KD})|} \quad (3.10)$$

$$fp(\mathcal{A}, \mathbb{KD}, \mathbb{UD}) = \frac{\sum_{UD_i} \begin{cases} 1 - \frac{1}{|\mathcal{A}(UD_i, \mathbb{KD})|} & \text{If } \mathcal{ID}(UD_i) \in \mathcal{A}(UD_i, \mathbb{KD}) \\ 0 & \text{Else if } |\mathcal{A}(UD_i, \mathbb{KD})| = 0 \\ 1 & \text{Else} \end{cases}}{|\mathbb{UD}|} \quad (3.11)$$

In the example of Table 3.1, we present an example of a re-identification attack for this policy with $\alpha = 0.25$. The attacker outputs two identities $\{U_{i_1}, U_{i_2}\}$. Let us say that the correct identity is U_{i_2} then $p = 0.5$, while it is $p \simeq 0.33$ for the top-k policy. And for the false positive rate, $fp = 0.5$ for the threshold-based and $fp = 0.66$ for the top-3. It is also worth mentioning that $fp \neq 1 - p$ since the precision considers that outputting no response gives a precision of 0 while the false-positive rate favors outputting no results rather than giving a set of false ones (i.e., $fp = 1$ for an empty response).

3.4 EVALUATION

We present in this section the evaluation of AP-Attack. We start by presenting the attacks and LPPMs used in this evaluation and how they have been configured (Section 3.4.1 and 3.4.2), our used datasets (Section 3.4.3) and our experimental setup (Section 3.4.4). Then, we present the performance of our proposed AP-Attack compared to state-of-the-art attacks (Section 3.4.5). We then demonstrate the lack of resilience of three representative LPPMs of the literature (Section 3.4.6). Both using the single-output policy. In Section 3.4.7, we use the top-k policy to evaluate the anonymity set size of users protected with different LPPMs. In Section 3.4.8, we show the advantages of the non-systematic threshold-based policy compared to a systematic policy such as top-k in assessing the effectiveness of LPPMs. Finally in Section 3.4.9, we analyze the effect on the re-identification of both the number of users in the system and the proportion of mobility available for each trace.

The evaluation answers the following questions:

- What is the most effective attack between POI-Attack, PIT-Attack and AP-Attack? (Section 3.4.5).
- What is the most effective LPPM between Geo-I, Promesse and W4M against the considered attacks? (Section 3.4.6).
- What it is the size of the anonymity set for each LPPM using the top-k policy? (Section 3.4.7).
- What is the most suitable policy for an attacker between the top-k and threshold-based policy? (Section 3.4.8).
- What is the impact of the number of users in the system and the proportion of the users' mobility available to the attacker? (Section 3.4.9)

3.4.1 Attack Competitors

In this section, we describe POI-Attack [82] and PIT-Attack [30] two state of the art attacks against which we compare the performance of AP-Attack.

Points Of Interest Attack - POI-Attack

This attack uses Points of interest (POIs) to characterize users' profiles. Therefore $\mathcal{P}_{poi}(KD_i)$ is the set of POIs extracted from the trace KD_i . Those points are extracted using clustering algorithms such as the ones presented in [107] [44] parameterized with the diameter of a geographical zone where a user has stopped and a minimum duration characterizing her stop. To measure the similarity between two sets of POIs, each POI of the first set is associated with the geographically closest POI in the second set. The dissimilarity between the two sets will be equal to the median of all the geographical distances, which is computed as presented in Equation 3.12. Where X and Y are the sets of POIs for each trace and $d(X_r, Y_t)$ computes the geographical distance between two POIs X_r and Y_t .

$$d_{POISets}(X, Y) = median \left[\left\{ \min_t [d(X_r, Y_t)] \setminus \forall r \right\} \cup \left\{ \min_r [d(X_r, Y_t)] \setminus \forall t \right\} \right] \quad (3.12)$$

Probabilistic Inter-POI Transition Attack - PIT-Attack [30]

In addition to extracting POIs, this attack takes into consideration the transition probability from one POI to another. Specifically, the authors rely on mobility Markov chains [28] where the states are POIs ($P = P_1, P_2, \dots, P_k$) ordered by the number points in each POI and the edges' labels are transitions probabilities between POIs (t_{P_i, P_j}). This is done by computing the proportion of transition between each POI in the mobility traces. In order to compute the distance between two mobility Markov chains, two pieces of information are taken into account: the geographical distance between POIs and the weight of each POI. The weight of a POI is computed using the proportion of points contained inside the POI. More precisely the authors proposed many distance metrics to compare Markov chains. The most effective one is the *stats-prox* distance which is a combination of two distances: the stationary distance and the proximity distance (Equation 3.13). The stationary distance (Equation 3.14) sums the weighted geographical distances between each combination of two POIs if the distance is lower than a parameter d_0 . And the proximity distance (Equation 3.15) after ranking the POIs by their weight in each Markov Chains. It adds scores r_i if two POIs of the rank i are closer than a parameter Δ . The score is halved after each rank $r_i = \frac{1}{2}r_{i-1}$ and r_0 is a parameter. The dissimilarity between the two Markov chain is the inverse of the total score.

$$d_{stats-prox} \equiv if(d_{stat} \leq \gamma \text{ and } d_{prox} \leq 10^5 km) \ d_{stat} \text{ else } d_{prox} \quad (3.13)$$

$$d_{stats}(P, Q) = \sum_{P_i, Q_i \in P \times Q} w(P_i) \times \begin{cases} d(P_i, Q_j) & \text{If } d(P_i, Q_j) < d_0 \\ 0 & \text{Else} \end{cases} \quad (3.14)$$

$$d_{prox}(P, Q) = \left(\sum_{i=1}^{\min(|P|, |Q|)} \begin{cases} r_i & \text{If } d(P_i, Q_i) < \Delta \\ 0 & \text{Else} \end{cases} \right)^{-1} \quad (3.15)$$

The above two attacks rely almost exclusively on POIs, eliminating the information contained inside the trajectories. Also, LPPMs focusing on the elimination of POIs yield to an ineffective attack as illustrated in Section 3.4.6.

3.4.2 Attacks and LPPMs Configuration

The three attacks evaluated in this chapter AP-Attack, POI-Attack and PIT-Attack have several configuration parameters. Specifically, AP-Attack has a cell size parameter that we have fixed at 800 meters in this evaluation. After a number of calibration experiments, we have chosen this value because it was big enough to include POIs and was resilient to noisy traces (for instance against GeoI or noisy GPS coordinates). In addition, re-identification rates result for cells between 50 meters and 800 meters are approximately similar. Furthermore, POI-Attack and PIT-Attack require parameters for the extraction of the POIs from the traces. These parameters are the diameter of the clustering area (that we fixed at 200 meters) and the minimum time spent inside a POI (that we fixed at 1 hour). These values have been chosen after a series of experimentations yielding to the best results. It is worth mentioning that in [82] POI-Attack was used in a different configuration. Indeed, the authors re-identified the obfuscated mobility traces against the non-obfuscated version of those traces, rather than using past mobility as a training knowledge. In consequence, re-identification is easier.

To evaluate AP-Attack, we have chosen three representative LPPMs of the literature: (1) Geo-I, which adds Laplacian noise to mobility traces and enforces a guarantee inspired from Differential privacy; (2) Promesse, which uses speed smoothing to erase POIs and (3) W4M, which alters traces to group them in cylindrical volumes hence enforcing k-anonymity. Each LPPM has several configuration parameters. These parameters have an impact on the privacy level offered to the users but also on the quality of the resulting obfuscated data. We have decided to configure each LPPM following a medium level of protection. This choice is motivated by the fact that our objective is not to find the best LPPM configuration but rather to show that with a reasonable alteration of the data, the LPPMs do not succeed completely at protecting the user from re-identification. Other experiments with other configurations of the used LPPMs or using other LPPMs of the literature can be done using our available toolkit [58]. Specifically, Geo-I is configured with

Table 3.2: *Description of datasets*

Name	CabSpotting	Geolife	MDC	PrivaMov
# users	536	42	144	48
Localization	San Francisco	Beijing	Geneva	Lyon
# records	11 219 955	1 574 338	904 422	973 684

a parameter ϵ that has an impact on the amount of noise added to the data (the lower epsilon the higher the noise). We have fixed the value of this parameter to 0.01, which corresponds to a medium privacy level. Promesse is configured with a parameter α that corresponds to the distance between two successive sampling points. We have fixed this parameter to 200 meters. Finally, W4M is configured with two parameters, k representing the minimum number of users inside the cylindrical volume and the radius δ of the latter. We have fixed these parameters at $k = 2$ and $\delta = 600$ meters because W4M erases a lot of points making the dataset almost empty and those parameters guarantee privacy and availability of the data.

3.4.3 Datasets

We used four real mobility datasets in our experiments. These datasets are: (1) Cabspotting [80] that contains the mobility of 536 cab drivers in the city of San Francisco; (2) Geolife [106] that contains the mobility of 42 users mainly in the city of Beijing; (3) MDC [54] that contains the mobility data of 144 users in the city of Geneva and (4) PrivaMov [9] that contains the mobility of 48 students and staff members in the city of Lyon. To make the comparison fair between the datasets, we selected in each dataset the 30 most active successive days. We present in Table 3.2 a description of the datasets used in our experiments. The users are not equally active in all the days of the period; some are more active than others. We consider as a mobility trace, the mobility of the user during all the period. In all the experiments described in this chapter, we split the datasets into a period of 15 days used for the training phase and 15 days used for the re-identification phase.

3.4.4 Experimental Setup

All of our experiments were carried out in a computer running an Ubuntu 14.04 OS with 50GB of RAM and 16 cores of 1.2Ghz each. Our testing application [58] written in Java & Scala and runs in the Java Virtual Machine 1.8.0.

3.4.5 Evaluation of Re-identification Attacks with Single Output Policy

The first experiment we did was intended to compare the three considered re-identification attacks by measuring their re-identification rate on non-obfuscated data of the four considered datasets. The results are depicted in Figure 3.3. From this figure, we observe that AP-Attack outperforms the two other attacks on all the considered datasets. This experiment shows that sending mobility data "anonymously" (e.g., by using anonymous communication protocols such as TOR [21]) to application providers is not sufficient to protect the privacy of users as an adversary using re-identification attacks is able to recognize from 45% to 79% of the users in the four datasets. It is thus necessary for end-users to rely on LPPMs to protect their data. From this experiment, we also notice that Cabspotting is the dataset where the users are the most intrinsically protected. This comes from the fact that cab drivers have similar mobility patterns (e.g., they regularly go to the airport, famous hotels, malls and taxi parking places). Instead, MDC, GeoLife and PrivaMov are related to users having different mobility habits, which makes them easier to re-identify.

3.4.6 LPPMs Effectiveness Against Re-identification Attacks with Single output Policy

In this experiment, we compare the performance of the three considered LPPMs, i.e., Geo-I, Promesse and W4M. Specifically, we evaluate the re-identification rate obtained by the three former attacks on data obfuscated using these three LPPMs. Figure 3.4 shows the results of this experiment. Besides the three LPPMs, we report the results obtained for non-obfuscated data, which we use as a baseline. At first glance, we observe the high level of privacy enforced by W4M in the PrivaMov dataset




Figure 3.3: Performance of re-identification attacks on single output policy

(11%) and by Promesse in the Cabspotting dataset (6%) against AP-Attack, which is the most successful attack. Nevertheless, these two LPPMs seem not to be sufficient to protect users in the GeoLife and MDC datasets where the re-identification rate reaches 48% and 36% for W4M and 68% and 46% for Promesse. We notice that the LPPMs that erase POIs as Promesse and W4M nullify the attack POI-Attack and PIT-Attack. For instance, in the Geolife dataset, PIT-Attack goes from 47% for the non-obfuscated data to 0% with Promesse, while AP-Attack goes from 79% to 68%. Finally, we observe that Geo-I is the least efficient LPPM against re-identification attacks in the four datasets. We also notice that Geo-I affects less AP-Attack compared to POI-Attack. Indeed, AP-Attack goes down on average with -3% while POI-Attack goes down by -15% . The noise added to the points by Geo-I rarely gets them out of a cell, while the clustering algorithms used to form POIs suffer more from the noise. Summarizing, this experiment allows us to draw the following conclusions: (1) there is no one-size-fits-all LPPM, as the resilience of an LPPM to re-identification attacks depends on the underlying data; (2) users of a given dataset are not all equal in front of re-identification attacks, as on the four datasets there exist users that are never re-identified even in the absence of protection mechanisms (e.g., 54% for the

best case with Cabspotting and 21% for the worst case with Geolife).

3.4.7 Evaluation of Re-identification attacks with Top-K Policy

In this section, we present the result evaluation of Top-k policy in Figure 3.5. Instead of only measuring the user re-identification rate for this policy, we search for each user, which level of k needs to be set in order to find him. The higher a k is, the higher is the user's protection. We used for this evaluation AP-Attack only, since according to the previous results it greatly outperforms the other attacks. From the results, we first notice the singularity of the Cabspotting dataset. With a median ranging from $k = 3$ for the non-obfuscated data to a median of $k = 78$ for W4M. Even when the data is not obfuscated, the users are fairly safe with a third quartile of $k = 47$. For the other datasets, the values are different. In GeoLife, the third quintile with Geo-I is at $k = 1$. For Promesse only $k = 2$ and a little higher with $k = 14.75$ for W4M. This shows a real threat to users. Further investigation on the singularity of Cabspotting is conducted in Section 3.4.9 to see if the difference in the number of users is the reason why the users are protected or not.

3.4.8 Evaluation of Re-identification Attacks with Threshold-based Policy

In this section, we show the result of the threshold-based policy in both average precision and average false-positive rate. We compare it to the theoretical bounds of the top-k policy. Indeed, since the top-k policy is systematic at taking k users, which limits its results on both the considered metrics. Hence, for a given k , the average precision cannot go beyond $\frac{1}{k}$ (upper bound) and the average false-positive rate cannot go below $1 - \frac{1}{k}$ (lower bound). These bounds are obtained in the best case where the correct identity is always part of the k users outputted by the policy. The threshold-based policy is different since it does not systematically take a certain number of users (it can even output no identity if the confidence is not high enough). As a consequence, it can manage a low false-positive rate and even a better average

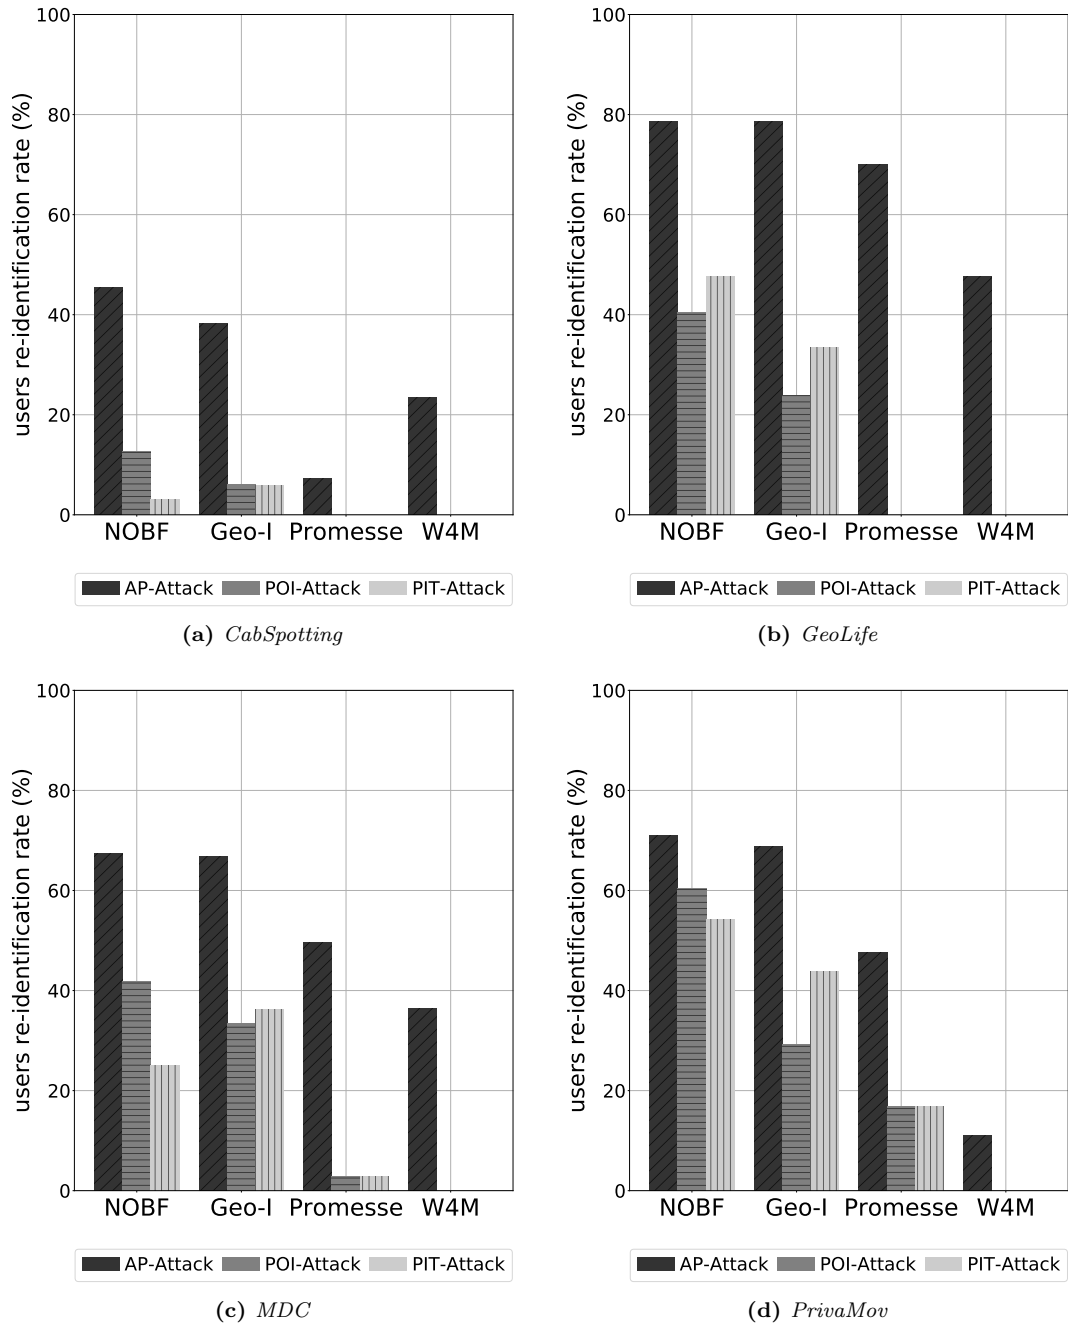


Figure 3.4: Performance of LPPMs against re-identification attacks (single output policy)

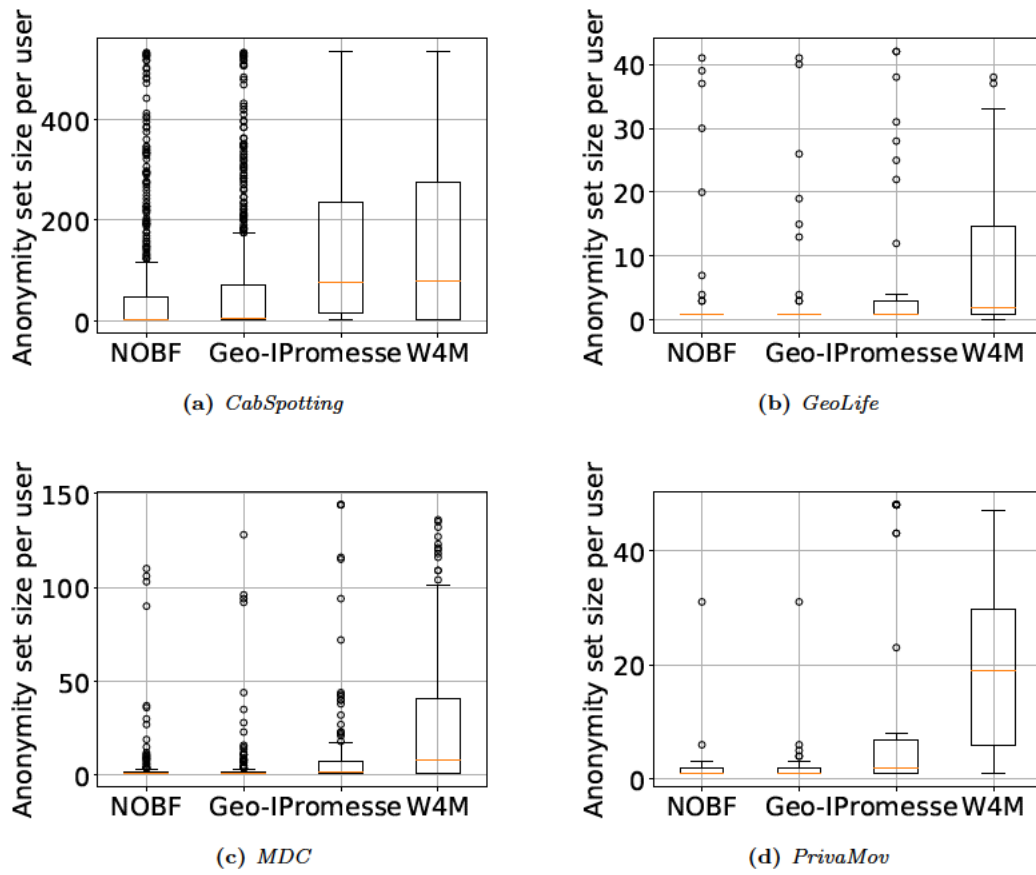


Figure 3.5: Performance of LPPMs

precision.

In Figure 3.6 the results for the GeoLife and Cabspotting dataset are depicted. We notice that for some values of α and some LPPMs, the threshold-based policy outperforms the theoretical bounds of the top-k policy in both the average precision and the average false-positive rate. This shows that taking a systematic number of users can lower the performance of the attack. We also notice that the best value of α is similar for each dataset.

It is worth mentioning that for the MDC dataset we have similar results except for $k = 2$ which has better performances. And for the Cabspotting dataset, the values of α need to be way lower (as a result of the close similarity between users) and the top-k method performs better in precision but lower in false-positive rate.


About the different LPPMs, it is interesting to notice that even after applying them the best results are obtained using the same α for one dataset. We also notice that *W4M* performs the best.

3.4.9 Analysis of the Parameters Affecting Re-identification


In this section, we analyze two parameters that can affect the results of the re-identification.

Number of Users Considered in the System


We present in Figure 3.7 the result of the user re-identification rate with respect to the number of users considered in the system. For this experience, we chose the dataset with the biggest number of users (536). We randomly picked a certain number of users. Then, we applied AP-Attack and computed the user re-identification rate (each result is the average of 5 random user picks). We notice that the number of users has an impact on the performance of the attack but not as high as expected. Indeed, starting from 200 users the rate stays steady. This comes from the fact that the addition of any new user can benefit - with regards to the protection against re-identification - only the users with similar behavior. This is why it is important for large scale systems to not only consider the total number of users but also the




(a) Cabspotting - Average Precision (higher the better)




(b) Cabspotting - Average False-Positive Rate (lower the better)




(c) Geolife - Average Precision (higher the better)




(d) Geolife - Average False-Positive Rate (lower the better)




(e) MDC - Average Precision (higher the better)



(f) MDC - Average False-Positive Rate (lower the better)



(g) PrivaMov - Average Precision (higher the better)



(h) PrivaMov - Average False-Positive Rate (lower the better)

Figure 3.6: Performance of the Threshold-Based Policy on Average Precision and Average False-Positive Rate Compared to the Theoretical Bound of the Top-k Policy




Figure 3.7: Effect of the number of users in the system to the user re-identification rate with single output policy (Cabspotting dataset)

similarity between them. In other words, if you have 1 billion users but in some areas few users are co-located then those users are at high risk. W4M is particular in this experience since the user re-identification rate increases at first. This is because W4M enforces a strong condition of the mobility trace to enforce k -anonymity and when it cannot enforce it the mobility traces are erased (which is considered non-re-identified but with an extremely bad utility). When there are few users the k -anonymity constraint is even harder to enforce. In consequence, more mobility traces are erased.

The Proportion of Trace Available

In this scenario, we consider an attacker that does not have access to the full user mobility trace but rather can eavesdrop on the users' mobility. We could imagine that the user uses various mobile applications and the attacker has access to the data of only a portion of vulnerable applications.

We want to analyze how much of the users' mobility an attacker needs to capture to be effective. To do that for each considered proportion of records p , we randomly




Figure 3.8: *Effect of the Proportion of the user mobility captured on the user re-identification rate with single output policy (Cabspotting datasets)*

select $p\%$ records for the traces of all the users and apply AP-Attack (the process is repeated 5 times). For the LPPMs considered, we apply the LPPM to the selected records rather than selecting from an already obfuscated trace. The results for the Cabspotting dataset are depicted in Figure 3.8. In this experiment, the absolute value of the rate is not the most important but we rather give importance to the relative value compared to the maximum obtainable value. Since the specificities of the dataset may bound how high the user re-identification rate can go. We notice in those results that with few records (around 20%) for the non-obfuscated and Geo-I protected version of the dataset, the attacker already obtains 80% of the best rate and already 64% for W4M. Promesse's results are particular since this method protects well the Cabspotting dataset. So every variation is not significant.

3.5 CONCLUSION

In this chapter, we presented a new model for re-identification attacks. It considers a two-step attacker that first uses profiling in order to compute for an anonymous

trace its probability to originate from a set of different identities. Then secondly, the attacker applies a policy in order to select a subset of identities to consider as probable identities. We also presented a re-identification attack based on a heat map representation of user profiles. We showed that this attack – which aggregates user mobility into a probability distribution acting as a fingerprint of user mobility – outperforms existing attacks on four real mobility datasets. Moreover, we studied the ability of three state-of-the-art LPPMs to protect users against re-identification attacks. The results showed that there is no one-size-fits-all LPPM. Instead, the degree of protection offered by LPPMs heavily depends on the underlying data. The use of the top-k policy showed that depending on the dataset and the protection method, the size of the anonymity set of the user can vary. Some users are more hidden thanks to their similar behavior with other users. The use of the threshold-based policy showed that a non-systematic method - that varies the number of identities taken into consideration - can be more beneficial to an attacker rather than systematically taking the same number of users.

In addition, in the chapter, we studied the effect of the number of users in the system. The results showed that after some increase in the number of users the re-identification rate stays steady. Mainly because the re-identification is not affected only by the number of users but also by the behavior of the added users taken into consideration. As a consequence, we would advice to not expect users to be protected only because they are part of a system with a big number of other users but rather to investigate the cluster of users and how the user hide themselves in groups (the top-k policy could be used for this). We also investigated the proportion of mobility an attacker needs to obtain in order to re-identify a user. The results show that with AP-Attack, a small portion of the user's mobility is enough for the attacker to reach its maximum potential of re-identification rate. Hence, this attack is able to build a proper user profile with few mobility data available (20% of randomly eavesdropped records for 80% of the attacks maximum potential).

In the next chapter, we propose a protection mechanism that is able to alter the profiles generated by the attacker and confuse the attacker on the owner of this profile. Since the experiments show that AP-Attack is the most effective attack, we are considering that heat maps better describe the users' behaviors and thus we use heat maps to guide our protection mechanism in order to transform (hence confuse) the attacker.

- Chapter 4 -

HMC: A Novel Location Privacy Protection Mechanism

Contents

4.1	Objectives and Roadmap	63
4.2	Recall on the Adversary Model	64
4.3	HMC Overview	65
4.4	Heat Map Alteration	66
4.5	Mobility Trace Reconstruction	68
4.6	Discussion on Alternatives for HMC	72
4.7	Experimental Evaluation of HMC	72
4.7.1	Privacy Metrics	72
4.7.2	Utility Metrics	74
4.7.3	Experimental Setup and Configurations	78
4.7.4	Privacy Evaluation	79
4.7.5	Utility Evaluation	85
4.7.6	Discussion	91
4.8	Conclusion	93

4.1 OBJECTIVES AND ROADMAP

As discussed in Chapter 2, various LPPMs have been proposed in the literature. They either enforce formal privacy guarantees (e.g., k-anonymity or differential privacy) or hide sensitive user information (e.g., Promesse hides POIs). In this chapter, we propose HMC (for *Heat Map Confusion*), a Location Privacy Protection Mechanism that protects users against re-identification attacks by reasoning on their mobility as a whole, captured using heat maps. Specifically, in order to protect a dataset of user mobility traces, HMC first extracts user profiles by aggregating the mobility of each user into a single heat map. Then, HMC alters each user heat map by making it look similar to the heat map of another user. To limit the decrease in data utility, HMC uses the heat map of the closest user as a basis for performing the alteration. Finally, HMC transforms back each altered heat map to a set of mobility traces by trying to retain as much as possible the users' original traces unchanged. The result is a protected mobility dataset on which an attacker that runs user re-identification attacks (e.g., AP-Attack, POI-Attack, PIT-Attack) fails in distinguishing between users.

In this chapter, the protection against re-identification attacks is evaluated with the single output policy and the anonymity set size of the top-k policy. Not only with one re-identification attack as in previous works, but with the results of multiple attacks. This allows us to demonstrate that HMC does not only protect the users against the attack that also uses heat maps to reason on user mobility but also against attacks that use other models (e.g., points of interest [82] or Mobility Markov chains [28]). Furthermore, we also evaluate data utility using multiple metrics that evaluate data distortion or the accuracy of applications.

To evaluate HMC we relied on four real mobility datasets (Cabspotting, Geolife, MDC, Privamov) and compared HMC with three representative adversaries (AP-Attack, POI-Attack, PIT-Attack). We also made HMC as an open-source prototype to reproduce our experiments (available at <https://github.com/mmaouche-insa/HMC>). The results show that HMC successfully decreases the user re-identification rate of all the attacks. Specifically, across all the datasets using HMC, 87% of mobile users are successfully protected against re-identification attacks, while others LPPMs only achieve a protection ranging from 43% to 79%. By considering only users protected

with a high utility, the proportion of users stays high for HMC with 75%, while for other LPPMs it goes down to proportions between 4% and 43%.

Roadmap In the remaining of this chapter, we start by giving a recall on the adversary model in Section 4.2. We present an overview of HMC in Section 4.3. Its two components heat map alteration and mobility trace reconstruction are described in Section 4.4 and Section 4.5 respectively. We discuss alternatives of HMC in Section 4.6. Experimental evaluation results are presented in Section 4.7. And finally, we draw our conclusions in Section 4.8.

4.2 RECALL ON THE ADVERSARY MODEL

We consider an attacker similar to the one presented in chapter 3. Let $\mathbb{U} = \{U_1, U_2, \dots, U_n\}$ be the set of users in the system and $\mathbb{KD} = \{T_1, T_2, \dots, T_n\}$ the set of background knowledge mobility traces previously gathered (T_i is the mobility trace of U_i). From each of these traces T_i , the adversary builds a user profile $P_i = \mathcal{P}(T_i)$ that characterizes the user mobility and acts as a fingerprint. Thus, the attacker has access to the set of profiles of the users in the system $\mathbb{P} = \{P_1, P_2, \dots, P_n\}$.

A Re-identification attack \mathcal{A} defined in Equation 5.3 run by the adversary using a single output policy tries to re-associate an anonymous trace T' from the Unknown user data \mathbb{UD} to a known user profile.

$$\begin{aligned} \mathcal{A} : \mathbb{UD} &\rightarrow \mathbb{U} \\ T' &\mapsto \mathcal{A}(T', \mathbb{KD}) = U_a \end{aligned} \quad (4.1)$$

Upon receiving an anonymous mobility trace T'_j , the adversary builds its profile $\mathcal{P}(T'_j)$ then researches in the background knowledge of profiles \mathbb{P} the most similar profiles with regard to a distance measure d and assigns its identity to the anonymous trace (See Equation 4.2).

$$\mathcal{ID}(T'_j) \leftarrow \arg \min_{U_k} d(\mathcal{P}(T'_j), P_k) \quad (4.2)$$

HMC Objective: Confuse the attacker so as $\arg \min_{U_k} d(\mathcal{P}(T'_j), P_k)$ is not the correct identity of T'_j .

4.3 HMC OVERVIEW

The process of obfuscating a mobility trace T whose identity $\mathcal{ID}(T) = a$ using HMC is depicted in Figure 4.1. This process is composed of three phases:

1. **Heat Map Creation (\mathcal{H}):** The objective is to construct the heat map of the mobility trace T waiting to be obfuscated. The method is based on the heat map representation of the mobility trace. In consequence, we start by computing $H = \mathcal{H}(T)$ using the heat map Construction module as done by AP-Attack.
2. **Heat Map Alteration (HMA):** The objective of this phase is to transform H into H' , an obfuscated heat map that is more similar to a user profile different than the one of user $\mathcal{ID}(T)$. There is actually more than one heat map that satisfies this property (See Equation 4.3), finding only one is sufficient.

$$\mathcal{HMA}(H, \mathbb{P}) = \{H' \mid \exists K : \mathcal{ID}(K) \neq \mathcal{ID}(H) \wedge \arg \min_{P_i \in \mathbb{P}} d(H', P_i) = K\} \quad (4.3)$$

3. **Mobility Trace Reconstruction (MTR):** We construct an obfuscated mobility trace T' whose heat map is H' the obfuscated heat map of H (Equation 4.4). We also use T to construct T' in order to keep the trace as similar as possible from the one before obfuscation with privacy guarantees as added value.

$$\mathcal{MTR}(H') = \{T' \mid \mathcal{H}(T') = H'\} \quad (4.4)$$

We describe in more details the two last phases (phase 1 is similar to the one of AP-Attack and described in Figure 3.2 of Section 3.2).




Figure 4.1: Overview of HMC

4.4 HEAT MAP ALTERATION

We need to construct H' a heat map that satisfies the property of the set $\mathcal{HMA}(H, \mathbb{P})$ defined in Equation 4.3. We chose to design a method based on iterative modifications. As depicted in Figure 4.2, we first search for U the most similar profile in \mathbb{P} and V the profile with the best utility (area coverage described in 4.7.2) in $\mathbb{P} \setminus \{U\}$.

$$U = \arg \min_{P_i \in \mathbb{P}} d(H, P_i) \quad (4.5)$$

$$V = \arg \max_{P_i \in \mathbb{P} \setminus \{U\}} \mathcal{UT}(H, P_i) \quad (4.6)$$

if $\mathcal{ID}(U) \neq \mathcal{ID}(H)$ then H already satisfies the property. This means that the user has a behavior (in the sense of the patterns of movements and the important locations) that is significantly different from her past mobility and does not need




Figure 4.2: Heat Map alteration iterative process

Algorithm 1 Algorithm of \mathcal{HMA} .

```

1: function  $\mathcal{HMA}(H, a, \mathbb{P}, n, IT_{max})$ 
2:    $U \leftarrow \arg \min_{P_i \in \mathbb{P}} d(H, P_i)$  The most similar profile
3:   if  $\mathcal{ID}(H) \neq \mathcal{ID}(U)$  then return  $H$  Does not need obfuscation
4:    $V \leftarrow \arg \max_{P_i \in \mathbb{P} \setminus \{U\}} \mathcal{AC}(H, P_i)$  Profile with the best utility
5:    $c \leftarrow 0$ 
6:   while  $d(H, V) > d(H, U) \wedge c \leq IT_{max}$  do
7:      $R \leftarrow n \cdot T$ 
8:      $W \leftarrow H \odot V \odot (1 - U)$   $\odot$  represents the pairwise product
9:      $O \leftarrow R + \left( \frac{a}{\sum W} \cdot W \right)$ 
10:     $H' \leftarrow \frac{1}{n} \cdot O$ 
11:    The counter  $c$  rewinds if  $H'$  gets closer to  $V$  compared to  $U$ 
12:     $c \leftarrow \text{update}(c, H, H', U, V)$ 
13:     $H \leftarrow H'$ 
14:  end while
15:  if  $c = IT_{max}$  then return  $V$  If no  $H'$  candidate is found, use  $V$ 
16:  return  $H$ 
17: end function

```

obfuscation (Line 3 of Algorithm 1). On the other hand, if the user is at risk of re-identification, the iterative process starts by searching H' .

We first transform the heat map back to a version with the number of records per cell rather than a frequency (Line 7). At each iteration a number of records are added to each cell, depending on the weigh computed using the formula in Equation 4.7. In order to affect as little as possible the \mathcal{UT} , we alter only cells that are already present in H . Furthermore, we want to reinforce points that are present in both H and V but that are not present in U . More specifically, in Algorithm 1 all the process of \mathcal{HMA} is presented. This algorithm stops after a number of iterations without improvement. In this case, V is used as H' since it satisfies the property of Equation 4.3 at the cost of utility loss.

$$\forall(i, j) : W_{ij} = H_{ij} V_{ij} (1 - U_{ij}) \quad (4.7)$$

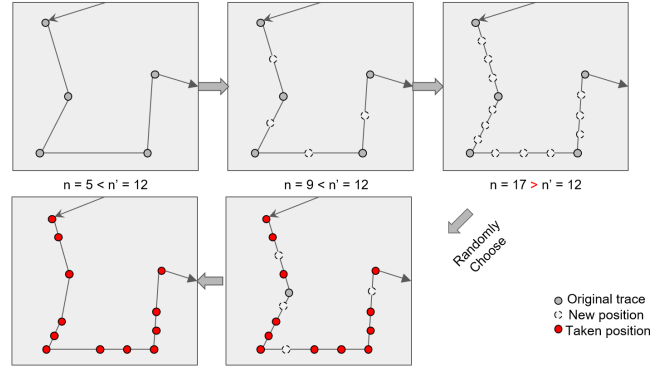


Figure 4.3: *Cell Number of Records Modification*

4.5 MOBILITY TRACE RECONSTRUCTION

This module generates T' a mobility trace whose aggregation is the heat map H' as expressed by Equation 4.4. The non-obfuscated mobility trace T is used in order to take into consideration utility metrics such as the spatial distortion \mathcal{SD} (Section 4.7.2). Even though, the abstraction using the heat maps loses the temporal aspects of the mobility traces, the original mobility trace is used to construct the protected one, in order to keep the temporal aspects as close as possible from the original trace.

We distinguish two types of cells in H' : (1) the cells that are present in T (ie., $H(i, j) \neq 0$) and (2) the cells that are not present in T . For the first case, Figure 4.3 illustrate how the traces contained inside a cell are altered in order to have the same intensity as the one in the obfuscated heat map H' . To reach this objective, we use a method inspired by the LPPM Promesse [83]. Specifically, as expressed in Algorithm 2, we use the interpolation between each pair of records in order to create new positions (the timestamp of a new position is equal to the center of the timestamps of the preceding and following record). We do this iteratively until we have enough positions as the number in H' (loop of line 3 to 11 of Algorithm 2). Lastly, we select randomly $H'(i, j)$ records and we keep the timestamps generated during the creation of the positions (line 12 of Algorithm 2). In the cases, where $H'(i, j) < H(i, j)$, we randomly select a set of records.

In addition to modifying the intensity of cells, HMC makes sure to not leave

small discriminating POIs. That's why, after transforming the number of records in a cell, we make sure to erase small size POIs.

Algorithm 2 Algorithm to adapt the number of records of \mathbb{R} to n records

```

1: function MODIFYNUMBEROFRECORDS( $\mathbb{R}, n$ )
2:    $\mathbb{P} \leftarrow \mathbb{R}$ 
3:   while  $|\mathbb{P}| < n$  do      Create new positions in the set  $\mathbb{P}$  until its size reaches  $n$ 
4:      $\mathbb{P}' \leftarrow ()$       Empty sequence
5:     for  $i \leftarrow 1$  to  $|\mathbb{P}|-1$  do
6:       Computing the latitude, longitude and timestamp of the middle point
7:        $p' \leftarrow (p[i-1] + p[i])/2$ 
8:        $\mathbb{P}' \leftarrow \text{appendToSequence}(\mathbb{P}', (p[i-1], p'))$ 
9:     end for
10:     $\mathbb{P} \leftarrow \text{appendToSequence}(\mathbb{P}, (p[|\mathbb{P}|-1]))$ 
11:  end while
12:  return  $\text{selectRandomly}(\mathbb{P}, n)$ 
13: end function

```

The second case happens only when no H' is found iteratively and V has to be used as a substitute. In this case, we need mobility data in those empty cells in order to apply the interpolation method described above. So, we use a set of records from the background knowledge in order to copy real mobility data. To be able to put new data, we use time gaps available in the trace (when the GPS is off for instance) to give temporal values to the records. Furthermore, we put a constraint on the portion of trace copied and the temporal gaps using a max speed limit v_{max} as illustrated in Figure 4.4. Where we have a trace with a time gap from the record a to the record b . As our objective is to generate realistic traces, we ensure that the selected interval is sufficient for a human to move (e.g., at least in walking speed and at most by car) from point a to the cell (i, j) then to b .

In Algorithm 3, we describe how an empty cell is filled with data. It uses as input: \mathbb{G} a list of all available time gaps in the trace, \mathbb{KD} a set of mobility traces to copy mobility from, (i, j) the coordinates of the cell to fill, v_{max} the maximum speed constraining the gaps as explained above, Δt_{max} that limits the time gaps inside a set of records (a set of records from one mobility trace is split into multiple sets of events that respect the limit Δt_{max}), θ_{limits} is the limit of the duration of the set of




Figure 4.4: *Time Gaps constraining method*

records to copy (to avoid copying a full day of mobility just to fill one cell).

To fill a cell, we first filter the time gaps according to v_{max} and (i, j) (Line 2) as depicted in Figure 4.4. Then, we assemble all the data available in the cell (i, j) from \mathbb{KD} after splitting it into multiple sets of records with respect to the constraint Δt_{max} . Next, from all the possible sets of records and all the possible gaps, we select the pair with least distance to connect one another, since a gap has a starting point and an ending point, the distance is the sum of the distances from the start of the gap to the set of records and from the set of records to the ending point of the gap.

Algorithm 3 Algorithm to fill an empty cell (i, j) with real mobility data from \mathbb{KD}

```

1: function FILLMOBILITYOFCELL( $\mathbb{G}, \mathbb{KD}, (i, j), v_{max}, \Delta t_{max}, \theta_{limit}$ )
2:    $(\mathbb{G}', p_{index}) \leftarrow filterGaps(\mathbb{G}, v_{max}, (i, j))$ 
3:   We keep only the gaps that verify the constrain of speed with respect to  $v_{max}$  and
    $(i, j)$  (see Figure 4.4)
4:    $possibleSetsOfEvents = \emptyset$ 
5:   for  $T$  in  $\mathbb{KD}$  with  $getEventsOfCell(T, (i, j))$  do
6:      $splittedSetsOfEvents \leftarrow splitEvents(getEventsOfCell(T, (i, j)), \Delta t_{max})$ 
7:     Split the events into multiple sets of events when the time gap between two
     records exceeds  $\Delta t_{max}$ 
8:      $possibleSetsOfEvents = possibleSetsOfEvents \cup splittedSetsOfEvents$ 
9:   end for
10:   $(setOfEvents, gap) \leftarrow bestMatch(possibleSetsOfEvents, \mathbb{G}, \theta_{limit})$ 
11:   $updateGaps(\mathbb{G}, gap)$  Either split the gap or erase it
12:   $out \leftarrow translateTime(setOfEvents, gap)$ 
13:  return  $out$ 
14: end function

```

4.6 DISCUSSION ON ALTERNATIVES FOR HMC

It has to be noted that both the method presented in Section 4.4 and Section 4.5 are pluggable with other methods. The only true conditions for HMC is to find both H' then T' that satisfies Equation 4.3 and Equation 4.4 respectively. In our instantiation of HMC, we use an iterative method to construct H' and in order to construct T' , we use a Time Distortion method [83] and a set of stored mobility traces to avoid using any outsourced library for synthetic mobility trace generation.

Fake user profiles can be used to transform the user behavior to distance her from the behavior of her past self. In this case, the fake profile generated needs to be close enough to the user to protect in order to maintain the data utility, but far enough to protect the user identity. With such a method, we gain security by avoiding the storage of real user profiles, but we lose the certainty that a user hides from its past self to look like a user that the attacker might re-identify.

4.7 EXPERIMENTAL EVALUATION OF HMC

In the following, we define the privacy metrics (Section 4.7.1) and utility metrics (Section 4.7.2) used in our experiments. In addition, we describe the experimental environment and configuration settings used in the experiments (Section 4.7.3). Finally, in our experiments, we compare the resilience of HMC to re-identification attacks with respect to state-of-the-art solutions in Section 4.7.4 and we further evaluate the utility of the data produced in Section 4.7.5. Our results show that across all the datasets, HMC outperforms its competitors in most cases. And for similar privacy results, HMC has better utility.

4.7.1 Privacy Metrics

We propose to evaluate the privacy levels offered by the LPPM, first with the evaluation of the effectiveness of the attacks using the re-identification rate of a single output policy, we also evaluate the anonymity size set of the attacks. In addition we

propose a multi-attack based privacy evaluation.

User Re-identification rate

We consider a single attack \mathcal{A} , a single set of background knowledge \mathbb{KD} and a set of testing mobility traces \mathbb{UD} . As a reminder, the user re-identification rate is a precision score of all the identities singly outputted by the attacker (see section 3.3.1) when re-identifying the traces of the set \mathbb{UD} . As described in Equation 4.8

$$r(\mathcal{A}, \mathbb{KD}, \mathbb{UD}) = \frac{\sum_{UD_i} \begin{cases} 1 & \text{If } \mathcal{A}(UD_i, \mathbb{KD}) = \mathcal{ID}(UD_i) \\ 0 & \text{Else} \end{cases}}{|\mathbb{UD}|} \quad (4.8)$$

k -Anonymity Set Metric:

In the adversary model, the attack outputs a single identity. For a user, Even though, not being designated as the most similar profile of her anonymized trace is a good news. Being the second or third most probable identity is still problematic. That is why we propose this k -anonymity metric. In order to measure for a certain tolerance level k , the proportion of users still at risk. This k represents the number of most probable identities for the anonymous trace being re-identified. More formally, the output of a k -attack $\mathcal{A}^{(k)}$ on an anonymous mobility trace T' is a set of k identities with the k most similar profiles. This privacy metric can be seen as a way to measure the k -anonymity set size of an obfuscated mobility trace.

Number of Successful Attacks:

This metric computes the number of successful attacks (ie., user correctly re-identified) on a user. It is defined in Equation 4.9 as a user-centric metric with $\mathbb{A} = \{\mathcal{A}_1, \mathcal{A}_2, \dots\}$ being the set of all attacks considered.

$$n(UD, \mathbb{KD}, \mathbb{A}) = \sum_{\mathcal{A}_k \in \mathbb{A}} \begin{cases} 1 & \text{If } \mathcal{A}_k(UD, \mathbb{KD}) = \mathcal{ID}(UD) \\ 0 & \text{Else} \end{cases} \quad (4.9)$$

Different methods of combining the attacks' results could be used (i.e., $\mathcal{A}' = f(\mathcal{A}_1, \mathcal{A}_2, \dots)$). For instance, we could leverage the rank results of all the attacks to choose as a result the profile with the best average ranking or use a voting system. Various tests were conducted but the results are inconclusive. Mainly, because AP-Attack is more efficient than the other two attacks and the cases where POI-Attack or PIT-Attack succeeds at re-identifying the correct user while AP-Attack fails are rare. In the end, this mix-up of attacks weakens AP-Attack. In consequence, we keep the multi-attack notion by counting the number of successful attacks but we mainly focus on finding the cases where no attack succeeds. This includes the strongest attack (in our case AP-Attack) but also the cases where POI-Attack or PIT-Attack are the only successful attacks.

4.7.2 Utility Metrics

The goal of an LPPM is to protect the users' privacy. Unfortunately, the alterations made by the LPPM to the mobility data cause a decrease in the data's utility. Moreover, studies such as [13] make the observation that there is a trade-off between privacy and utility. In consequence, when designing an LPPM, it is important to evaluate the utility of the data produced. Indeed, designing a powerful LPPM that ensures users' privacy without considering the usefulness of the resulting data for later analysis is fruitless.

Two approaches arise when evaluating the utility of the altered data. The first is **data-centric**, which is generic and agnostic of the application. In this case, we consider that every application that is affected by the precision of the data is concerned and could profit from this metric. The second one is **application-centric**. In this case, we consider a particular application and the conclusion can only be generalized to applications with the same purpose.

In the remaining of this section, we describe the utility metrics used accompanied with examples of applications.

Area Coverage:

This metric computes how much the alteration affected the regions visited by a user [83]. In other words, while removing records makes places less significant for a user mobility (e.g., Erasing POIs), keeping the information on which regions the user goes through can be important. On the contrary, adding/moving records to new regions adds a piece of fake information that can lead to false deductions from the data analysis. For instance, we could conclude wrongly that a minor place has a large number of users going through it, which may push for ill-advised public investment in transport.

To compute the Area Coverage \mathcal{AC} , the map is divided into equal square regions. For T a mobility trace, $\mathcal{C}(T)$ (Eq. 4.10) returns the set of regions the user goes through, \mathbb{C} represents the set of all possible regions of the dataset and $e \sqsubset c$ means that the record e is inside the cell c

$$\mathcal{C}(T) = \{c \in \mathbb{C} \mid \exists e \in T : e \sqsubset c\} \quad (4.10)$$

To measure \mathcal{AC} of the obfuscation of T to T' , we compute the F-Score value of the precision-recall pair. The precision evaluates the proportion of cells the user goes through in the obfuscated trace that are present in the non-obfuscated trace. While the recall evaluates the proportion of cells of the non-obfuscated trace that are still found in the obfuscated trace.

An example of a use case could be the public health department searching for the areas in the city where noise disturbance is the most problematic by running a crowd-sensing campaign of noise levels in the city. Precise locations are not critical but covering the correct regions of the city is important.

$$\mathcal{AC}_{Precision}(T, T') = \frac{|\mathcal{C}(T) \cap \mathcal{C}(T')|}{|\mathcal{C}(T')|} \quad (4.11)$$

$$\mathcal{AC}_{Recall}(T, T') = \frac{|\mathcal{C}(T) \cap \mathcal{C}(T')|}{|\mathcal{C}(T)|} \quad (4.12)$$

$$\mathcal{AC}(T, T') = \mathcal{AC}_{F-Score}(T, T') = \frac{2 \cdot \mathcal{AC}_{Precision}(T, T') \cdot \mathcal{AC}_{Recall}(T, T')}{\mathcal{AC}_{Precision}(T, T') + \mathcal{AC}_{Recall}(T, T')} \quad (4.13)$$

Spatial Distortion:

This parameterless metric computes the spatial error. It considers the traces as polylines $T = (r_1, r_2, \dots)$ and $T' = (r'_1, r'_2, \dots)$. For each record x in T' we search for the minimal projection on T . $\mathcal{SD}(T, T')$ is the average of the minimal projection of all the records in T' .

$$\mathcal{SD}(T, T') = \frac{1}{|T'|} \sum_{x \in T'} \min_{0 \leq i < |T|} d_{\text{projection}}(x, r_i r_{i+1}) \quad (4.14)$$

An example of use case could be a city planner wanting to analyze the roads that need the most care by counting the number of users going through them. In this case, a precise spatial location to recognize the correct routes is essential.

Spatio-Temporal Distortion:

This metric computes a spatial error constrained by the timestamps of the records. As defined in Equation 4.16, the spatio-temporal distortion \mathcal{STD} is the average distance between each record of T' and its temporal projection into T . With, the temporal projection of the record $x = (x^{\text{lat}}, x^{\text{lon}}, x^t)$ in T' being its expected position r_e in T at time x^t . Specifically, we search for $r_i = (r_i^{\text{lat}}, r_i^{\text{lon}}, r_i^t)$ and $r_{i+1} = (r_{i+1}^{\text{lat}}, r_{i+1}^{\text{lon}}, r_{i+1}^t)$ in T such as $r_i^t \leq x^t \leq r_{i+1}^t$, then compute r_e the interpolation with the ratio $(x^t - r_i^t)/(r_{i+1}^t - r_i^t)$ (see Equation 4.15).

An example of use case could be, analyzing users' habits during the day. Such as, which places are mostly visited during the night and need more care in road lights.

$$\text{temporal_projection}(x, T) = \begin{cases} r_1 & \text{If } x^t < r_1^t \\ r_i + \frac{x^t - r_i^t}{r_{i+1}^t - r_i^t} (r_{i+1} - r_i) & \text{If } \exists i : r_i^t \leq x^t \leq r_{i+1}^t \\ r_{|T|} & \text{If } x^t > r_{|T|}^t \end{cases} \quad (4.15)$$

$$\mathcal{STD}(T, T') = \frac{1}{|T'|} \sum_{x \in T'} d_{\text{temporal_projection}}(x, T) \quad (4.16)$$

Distortion in Surrounding POIs:

This metric simulates an application that analyses the POIs surrounding the user location during her mobility. *Open Street Map* [62] is used for this metric. Their open data is uploaded to a MangoDB server and for each record x of the mobility trace in the obfuscated trace T' we query for the surrounding POIs in a rectangular area of size β (with $\mathcal{POI}(x, T, \beta)$). Then, we compare it to the result of the same query for the temporal projection of x in T (see Equation 4.15) using the harmonic mean of recall/precision (see Equation 4.17 & 4.18). The overall distortion in surrounding POIs is the average of all the F-scores of the records of T' (see Equation 4.19).

$$\mathcal{POI}_{Precision}(x, T, \beta) = \frac{|\mathcal{POI}(\text{temporal_projection}(x, T), \beta) \cap \mathcal{POI}(x, \beta)|}{|\mathcal{POI}(x, \beta)|} \quad (4.17)$$

$$\mathcal{POI}_{Recall}(x, T, \beta) = \frac{|\mathcal{POI}(\text{temporal_projection}(x, T), \beta) \cap \mathcal{POI}(x, \beta)|}{|\mathcal{POI}(\text{temporal_projection}(x, T), \beta)|} \quad (4.18)$$

$$\mathcal{DSP}(T, T', \beta) = \frac{1}{|T'|} \sum_{x \in T'} \frac{2 \cdot \mathcal{POI}_{Precision}(x, T, \beta) \cdot \mathcal{POI}_{Recall}(x, T, \beta)}{\mathcal{POI}_{Precision}(x, T, \beta) + \mathcal{POI}_{Recall}(x, T, \beta)} \quad (4.19)$$

This metric only evaluates if similar POIs are found. It can be extended further to a semantic metric by choosing only certain types of POIs while querying Open Street Map, using the "amenity" [61] categorization of the data that references the type of POI. For instance, one can search for sustenance POIs (i.e., bar, fast food, restaurant, cafe...) or for healthcare POIs (i.e., clinic, dentist, hospital, pharmacy...).

Number of Visits Distortion:

This metric simulates a data analysis where the number of visits to a place x is computed for a user. A visit is a record r_i that is within a radius α of x while r_{i-1} is not (See Eq.4.20). We compute the distortion between the number of visits in the obfuscated trace compared to the non-obfuscated trace (Eq.4.21).

$$\mathcal{NV}(T, x, \alpha) = |\{r_i \in T \mid d(r_i, x) \leq \alpha \wedge d(r_{i-1}, x) > \alpha \wedge 1 < i \leq |T|\}| \quad (4.20)$$

$$\mathcal{NV}\mathcal{D}(T, T', x, \alpha) = \frac{|\mathcal{NV}(T, x, \alpha) - \mathcal{NV}(T', x, \alpha)|}{\mathcal{NV}(T, x, \alpha)} \quad (4.21)$$

4.7.3 Experimental Setup and Configurations

The following experiments were conducted in a computer running an Ubuntu 14.04 OS with 50GB of RAM and 16 cores of 1.2Ghz each. The HMC prototype is developed in Java & Scala and runs in the Java Virtual Machine 1.8.0. It is available for download at: <https://github.com/mmaouche-insa/HMC>

In our experiments, we compare HMC with three state-of-the-art LPPMs: Geo-I, Promesse and W4M. The LPPMs come with their own configuration parameters, that are set as follows. Geo-I's ϵ configuration parameter is set to 0.01; this adds a medium amount of noise to the obfuscated data (the lower ϵ the higher the noise). Promesse's α configuration parameter is set to 200 meters, it represents the distance between two successive sampling points. W4M has two configuration parameters, i.e., k that is the minimum number of users inside the cylindrical volume, and δ the radius of the cylindrical volume. Here, k and δ were respectively set to medium values 2 and 600 meters. Finally, HMC's cell size is set to 800 meters (similar to the good configuration of a heat map based attack)

Furthermore, to stress the robustness of the LPPMs and thus evaluating the privacy level they provide, we consider three re-identification attacks in our experiments, namely PIT-Attack, POI-Attack and AP-Attack. The implementations of these attacks have their own configuration parameters. PIT-Attack and POI-Attack have two parameters for the extraction of the POIs from the traces. These parameters are the diameter of the clustering area, and the minimum time spent inside a POI. They were respectively set to 200 meters and 1 hour. And AP-Attack has a configuration parameter that corresponds to the cell size, and that was set to 800 meters. Finally, to evaluate the data utility level provided by the LPPMs, we consider the three utility metrics (described in Section 4.7.2) that are configured as follows. The Area Coverage utility metric has a configuration parameter that represents the size of a square region, it is set to 800 meters. For the metric evaluating the F-score of the surrounding POIs. Its square bounding-box is of distance 200 meters from the record




Figure 4.5: Comparison of HMC with competitors - Robustness against multiple attacks - CabSpotting & GeoLife datasets

considered. The utility metric that corresponds to the Number of Visits Distortion has one parameter α set to 100 meters, which is the distance threshold α from the place to the record considered as a visit. And the spatial and spatio-temporal distortion utility metrics do not need configuration.

4.7.4 Privacy Evaluation

In this section, we compare HMC against three LPPMs using three re-identification attacks.

Resilience against Multiple-Attacks:

Let us start with the proposed multi-privacy metric. Indeed, we merge the numerous attacks as presented in Equation 4.9 of Section 4.7.1. In Figures 4.5 and 4.6 we present the results by showing the proportion of users with their corresponding number of successful attacks. We notice that HMC behaves well with a 0 attack protection of 65% to 94% while W4M does 21% to 89% and Promesse 29% to 92%.

If we compare the proportion per dataset (i.e., with $proportion_{LPPM}$ versus $proportion_{HMC}$), HMC has a proportion of users of -16% better than W4M and




Figure 4.6: Comparison of HMC with competitors - Robustness against multiple attacks - MDC & PrivaMov datasets

—22% better than Promesse on average. If we consider all the users (across all the datasets), 87% of the users obfuscated with HMC have 0 successful attacks. While it is 78% for Promesse, 72% for W4M and 48% for Geo-I. This shows that HMC outperforms the other competitors, even when it is facing two out of three attacks that use a different underlying model (i.e, POI-based attacks)

Anonymity Set:

For the anonymity set size experiments, the results are depicted in Figure 4.14. We notice that for $k = 2$, even though, the gap is getting smaller HMC still outperforms the other LPPMs in three out of the four datasets. The confusion method of HMC does not transform the mobility to make it the second most similar to its past self but rather to look similar to another user. This is why the correct user does not fall off to the second position. But as depicted in the results, we notice that until $k = 5$ users are better hidden with HMC. This result comes from the fact that the target profile for the confusion is the one with the best utility in area coverage (the confusion is limited to a profile with low utility distortion). Hence, better anonymity size results can be obtained with HMC by selecting other types of target users (e.g., randomly or k -st most similar profiles) but with the cost of lowering the utility. As, it is shown in the next utility experiment, this configuration of HMC is capable of providing good privacy protection with a better utility than its competitors.

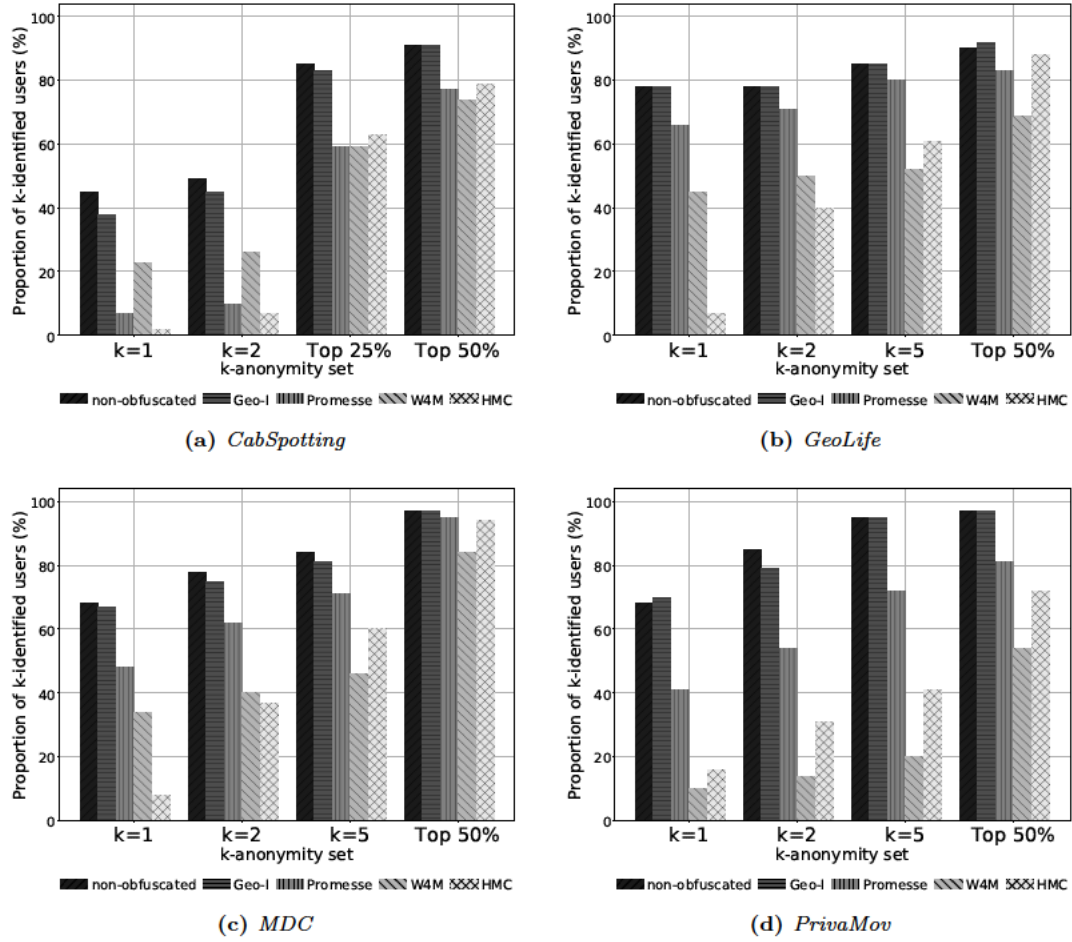


Figure 4.7: Detailed comparison of HMC with competitors - Anonymity set size against AP-Attack

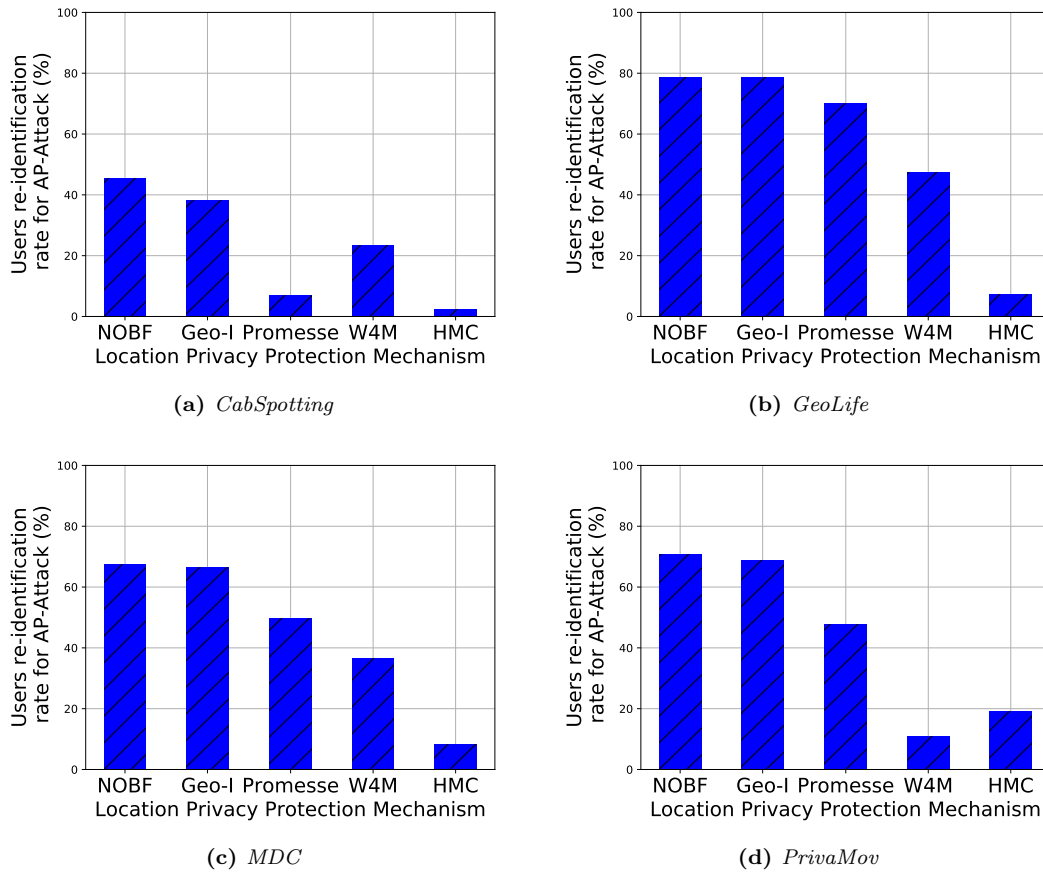


Figure 4.8: Detailed comparison of HMC with competitors - Robustness against AP-Attack

Detailed Resilience against Each Individual Attack:

Figures 4.8, 4.9 and 4.10 present the detailed results of user re-identification rate per type of attack, for respectively, AP-Attack, POI-Attack and PIT-Attack. We first notice that against the strongest attack AP-Attack, HMC behaves the best. In 3 out 4 of the datasets the rate ranges from 2% to 8% while W4M's rates range from 23% to 48%. In the other dataset PrivaMov, W4M performs better with 11% over the 19% of HMC. On average HMC has -20% of user re-identification rate (ie., $r_{W4M} - r_{HMC}$). For POI-Attack and PIT-Attack, HMC performs worse than W4M but still has low re-identification rates $< 20\%$.

In conclusion, HMC outperforms the other LPPMs vastly on AP-Attack which was expected since HMC is based on the heat map representation of the users'

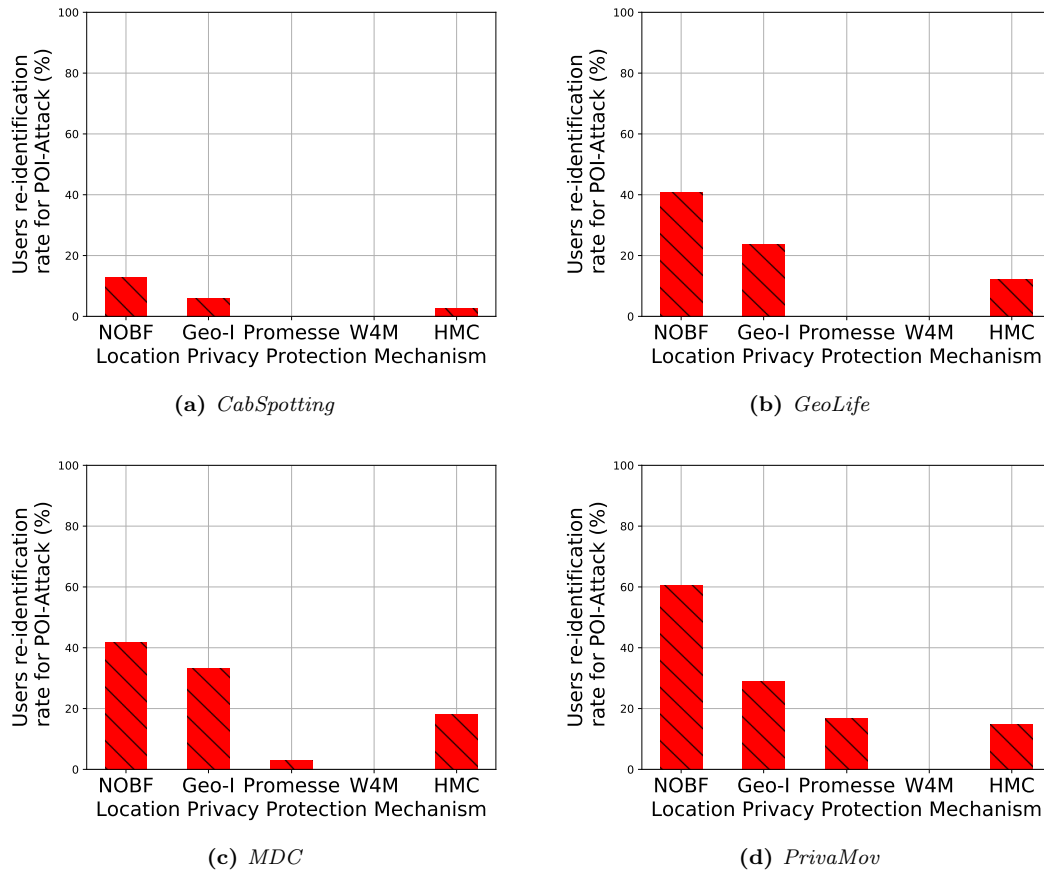


Figure 4.9: Detailed comparison of HMC with competitors - Robustness against POI-Attack




Figure 4.10: Detailed comparison of HMC with competitors - Robustness against PIT-Attack

Table 4.1: *Utility Measure Levels Description*

	\mathcal{AC}	\mathcal{SD}
Low	≤ 0.8	$> 200meters$
High	> 0.8	$\leq 200meters$

mobility. While having good performing results on attacks based on POIs.

4.7.5 Utility Evaluation

In this section, we present the utility results of HMC in area coverage, spatial distortion, spatio-temporal distortion, the distortion in surrounding POIs and the distortion in the number of visits.

Data-Centric Utility:

To present clearly the results, all the metrics have a threshold value in which the utility becomes too low for the user. The Table 4.1 presents those thresholds and the results are depicted in Figure 4.11, only the results for the users fully protected by the LPPM are presented (ie., 0 successful attacks) because measuring the utility of a non-protected user is not meaningful for an LPPM.

We notice that HMC has a big portion of users with High \mathcal{AC} and High \mathcal{SD} ranging from 27% to 89%, while these metrics for ange from 2% to 5% and Promesse 4% to 35%. If we consider all the users across all the datasets, 75% of the users that use HMC are fully protected against re-identification attacks and have a high Area Coverage and Spatial Distortion. While it is only 43% for GeoI, 27% for Promesse and as few as 4% for W4M. Overall the only datasets where HMC is challenged in terms of privacy is by over-altering the data and thus lowering the utility. This is the case for PrivaMov where W4M has better privacy but few users have a high utility (only 2%). In Geolife also, Promesse has comparable privacy results but overall half of the users are protected at the cost of lower utility while HMC protects most of them with high utility.

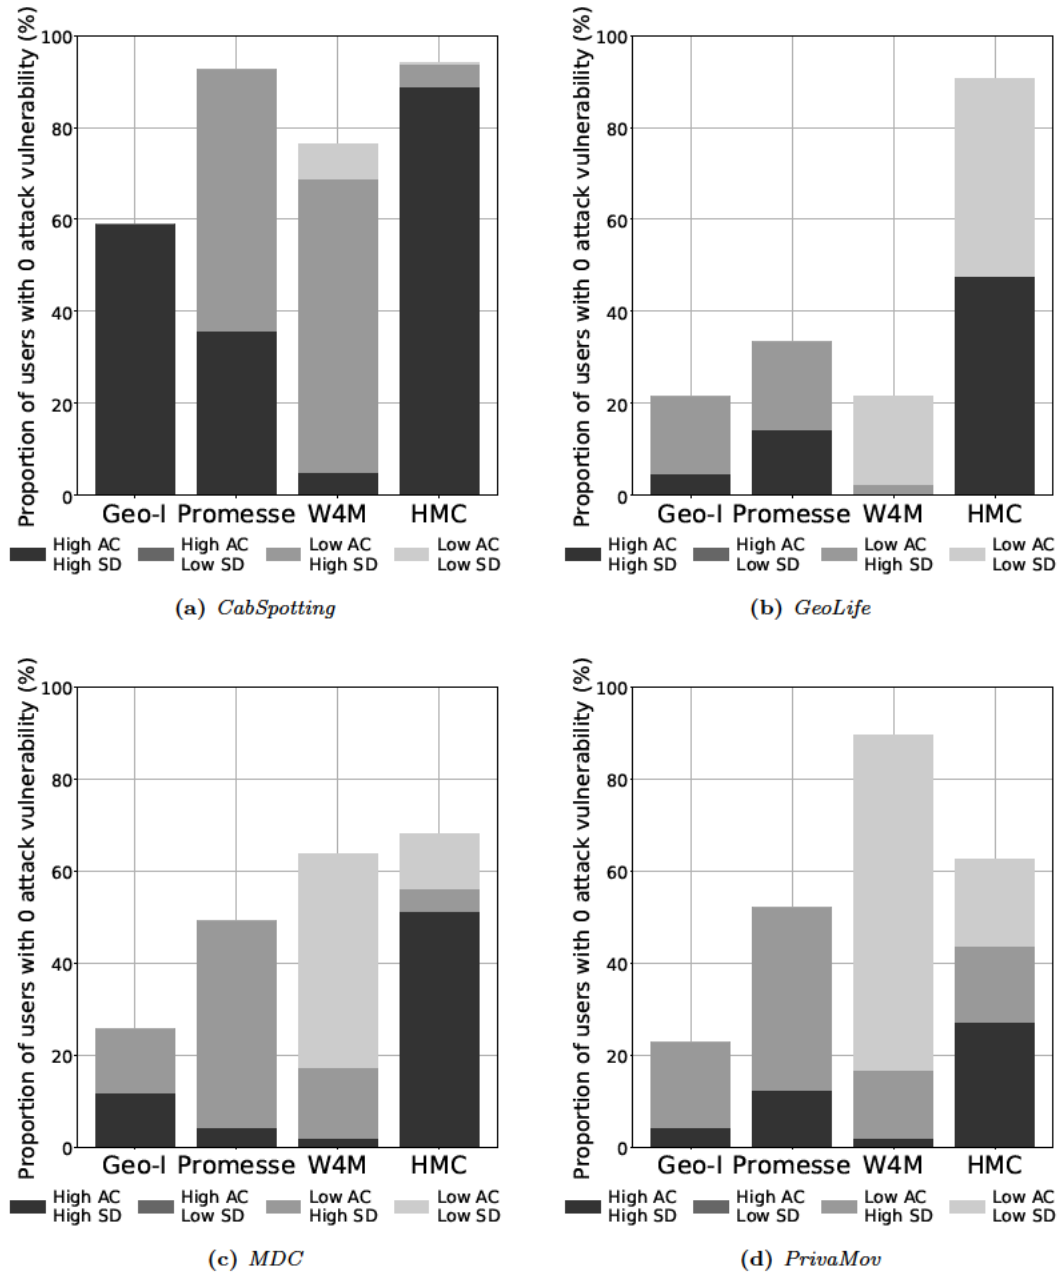


Figure 4.11: Detailed comparison of HMC with competitors - Multi-utility metrics evaluation

A more detailed analysis for Area Coverage is depicted in Figure 4.12. We notice that HMC outperforms all the other LPPMs in terms of Area Coverage. HMC's F-score average ranges from 0.63 to 0.98 while W4M's average ranges from 0.15 to 0.68, for Promesse it is from 0.53 to 0.75 comparable to HMC but still lower in each dataset.

In term of Spatial Distortion, we present the separate result in Figure 4.13. HMC's distortions are very low, each median is in centimeters in the four datasets. For W4M, the medians of its distortions range between $0m$ to $3.6Km$. HMC has lower values (excluding extreme cases) thanks to the Promesse-like interpolation technique that creates low spatial distortion because it puts new records only of the trajectory of the original trace. This is why Promesse has medians ranging from 4 meters to 13 meters.

In terms of spatio-temporal distortion, the results are presented in Figure 4.14. We first notice the results are worse than the spatial distortion. Indeed, the spatio-temporal distortion is the constrained version of the later. Promesse is the LPPM that suffers the most from the temporal constraint, as this method uses time distortion in a large portion of trace with speed smoothing in order to erase POIs. Across all the datasets, There are 76% of the users protected with a spatio-temporal distortion greater than 200 meters for Promesse. In contrast, there is only 27% of users for HMC. Also, the proportion of users protected with a spatio-temporal distortion lower than 10 meters is of 50%. W4M already had bad results for the spatial distortion, with a more constraining metric, there is 71% of users across all the datasets that are protected but with a spatio-temporal distortion greater than 200 meters. For Geo-I, even though only 50% are fully protected, there is a systematic noise added to the records, so there is always a distortion around 200 meters.

Application-centric Utility:

We present the result of the comparison of HMC to the other LPPMs with the utility metric that measures the F-score of the query of surrounding POIs (section 4.7.2) in Figure 4.15. We first notice that with the configuration of 200 meters for the rectangular area size, the average F-score is quite low. W4M performs better in Cabsptotting for $[0.75, 1]$ interval but on average since HMC has 64% of users in the

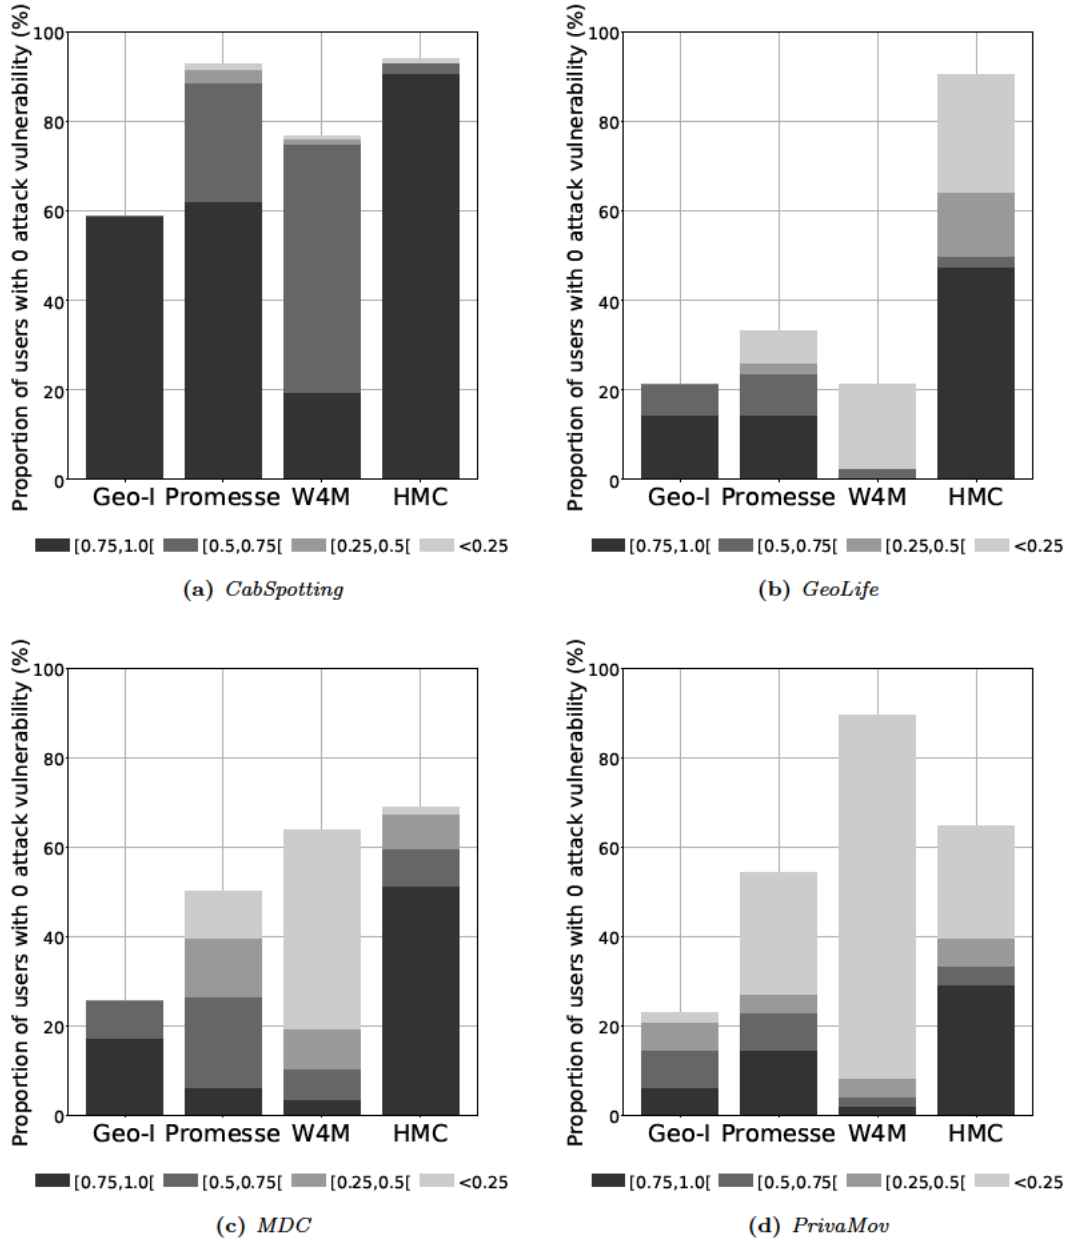


Figure 4.12: Detailed comparison of HMC with competitors - Area coverage utility metric

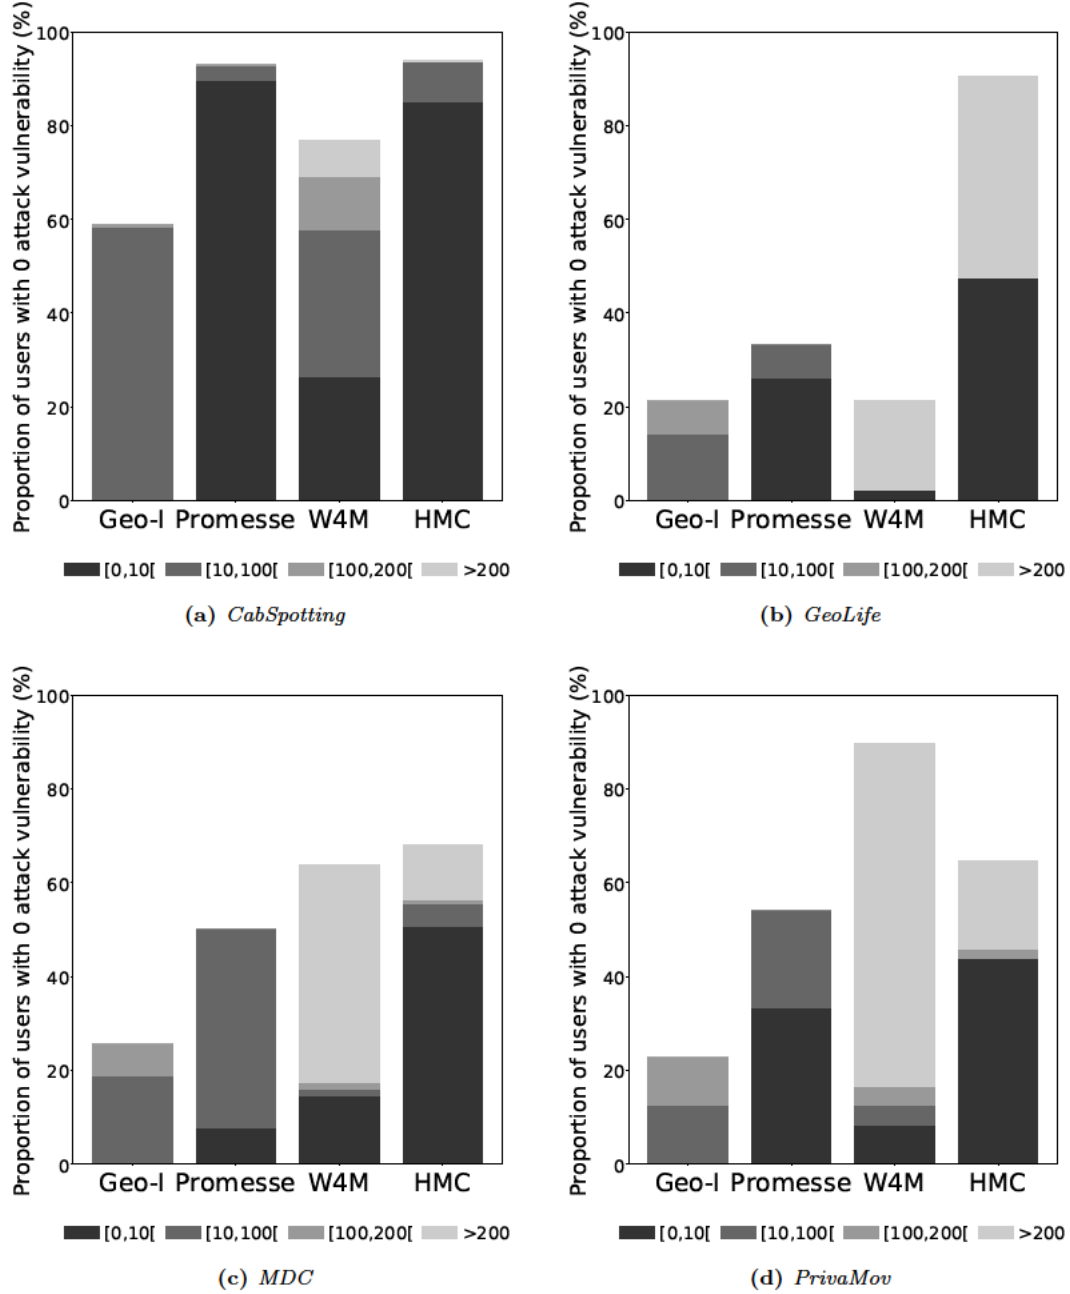


Figure 4.13: Detailed comparison of HMC with competitors - Spatial distortion utility metric

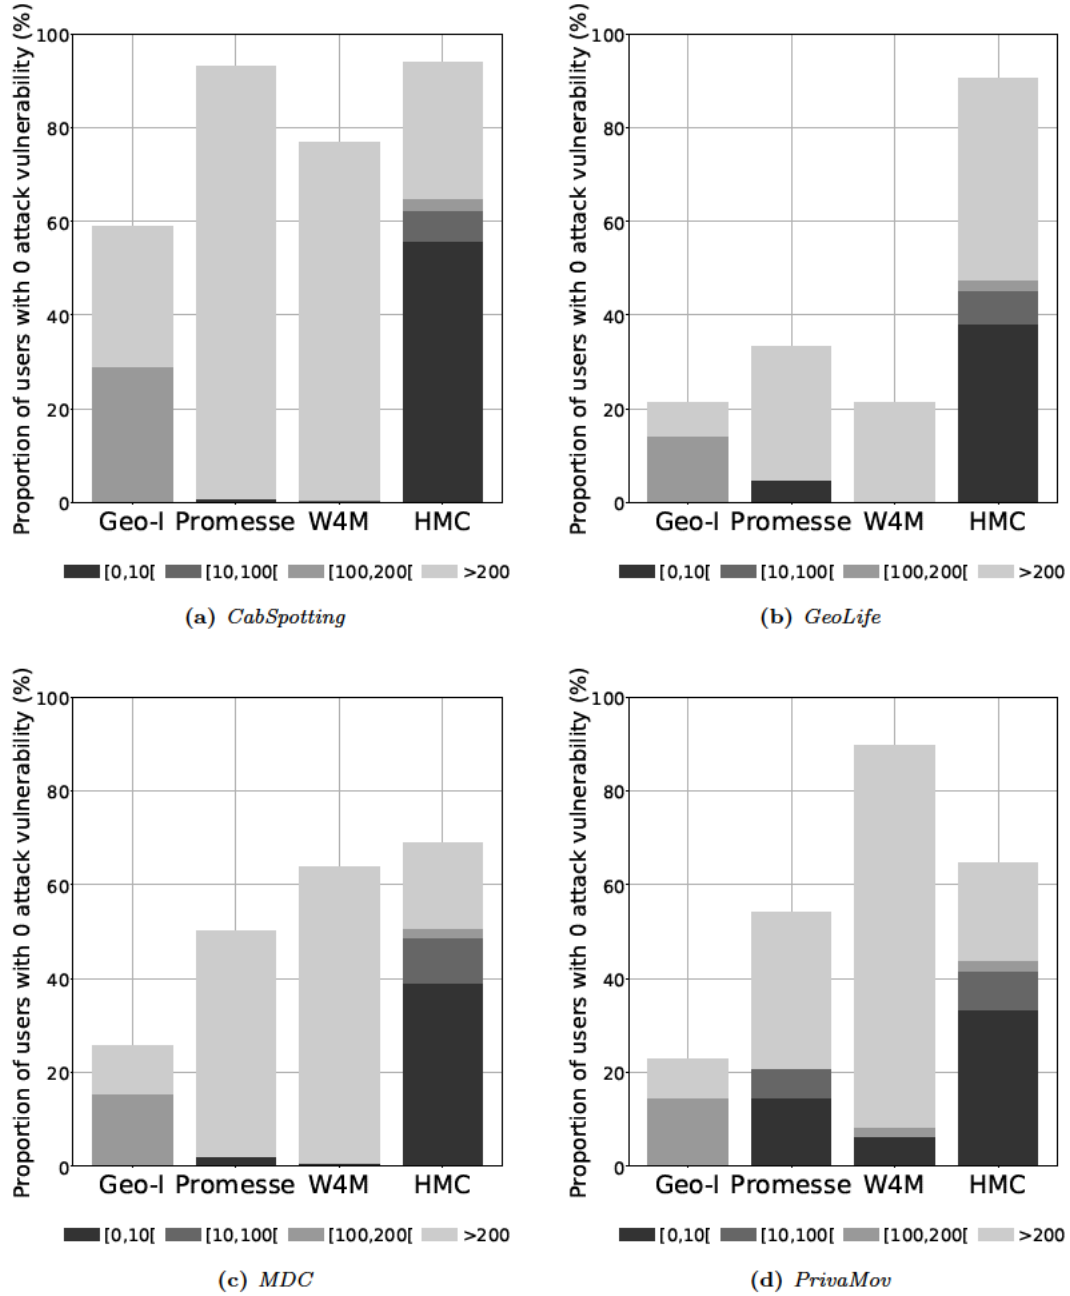


Figure 4.14: Detailed comparison of HMC with competitors - Spatial-temporal distortion utility metric

[0.5, 0.75[interval its average F-score is better (0.37 compared to the average F-score of 0.30 of W4M). Except for Promesse whose average F-scores by dataset ranges from 0.1 to 0.12 the other LPPMs have similar results with a small lead for HMC. Since, HMC F-scores ranges from 0.13 to 0.39, for W4M it is from 0.13 to 0.30 and Geo-I from 0.11 to 0.42.

For the last utility experiment, we present the result of the visits of "Union Square" in San Francisco (CabSpotting Dataset). We first notice the good results of Promesse by construction with 90% out of the 92% fully protected users have a distortion lower than 0.25. HMC has similar good results with 81% out of the 94% fully protected user with a distortion lower than 0.25. W4M has a diversity of users with two 30% groups of users with respectively 0.25 to 0.5 distortion and 0.5 to 0.75 distortion, this another low utility level for W4M.

4.7.6 Discussion

HMC has good results in utility because it aims at altering the data as few as possible. The cases where new cells of the map are filled are rare and those are the cases where the utility is deteriorated.

We notice that Promesse has lower utility results, not because of its perturbation method (which is utility-preserving) but rather because it does not manage well big time gaps where the user movement was not recorded. In those cases, Promesse fills those gaps with mobility data that adds distortion to the data. Also, Promesse and Geo-I apply a systematic perturbation method, even if the user does not need much altering in order to be protected, the utility is always lowered (but still the best to erase POIs). Most importantly, while utility-wise, it has good performances, Promesse's poor privacy-results (particularly against AP-Attack) makes it a bad candidate to protect against the user re-identification threat.

On the other hand, W4M performs poorly utility-wise even in the Cabspotting dataset where numerous users and records are available. Its results on POI-based attacks are good. but far from convincing against AP-Attack. This actually as stated before, the motivation behind the design of the heat map based protection mechanism HMC.

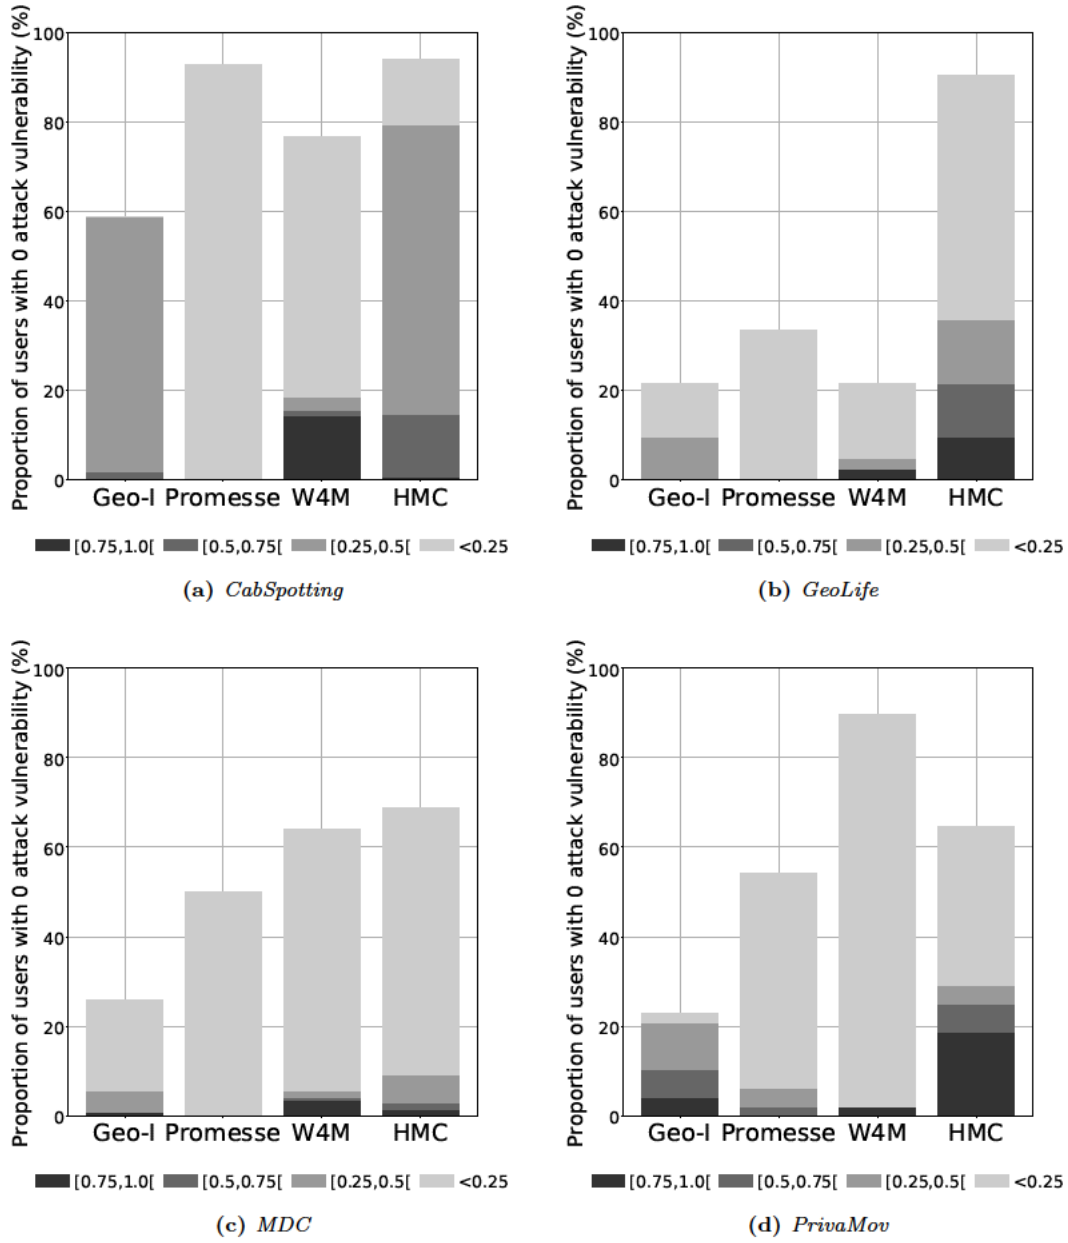


Figure 4.15: Detailed comparison of HMC with competitors - F -scores of surrounding POIs query utility metric




Figure 4.16: Comparison of HMC with competitors - Utility metric in terms of users' number of visits distortion – Cabspotting dataset

For the case of Geo-I, adding noise deteriorates the utility more than Promesse but it is inept to protect against re-identification attacks, having results similar to non-obfuscated traces. This is due to the dependency between successive records. Indeed, this makes the ϵ -GeoI guarantee loses its power to a $n\epsilon$ -GeoI (n being the number of records). In addition, practically Geo-I guarantee is on location attacks but it still permits to locate users in blocks (depending on ϵ), which is enough for AP-Attack. In order to affect the heat map representation of the user profile, Geo-I would need to add significant noise.

4.8 CONCLUSION

In this chapter, we presented HMC a novel LPPM that protects users against re-identification attacks. It uses a heat map alteration process in order to confuse the attacker and to make the re-identification fall to the wrong user. The solution proposed to implement HMC is based on an iterative modification to transform the heat map and an interpolation technique to alter the number of records in the

mobility trace. The heat map is a good abstraction of the mobility as it takes into consideration higher level features that can discriminate between users.

HMC was evaluated on four real mobility datasets against three representative re-identification attacks and compared to three competitive LPPMs. The evaluation was done using a multi-privacy metric which computes the number of successful re-identification attacks and a multi-utility metric with a threshold-based Low/High utility categorization simple to interpret. The results show that HMC outperforms the other LPPMs in terms of both privacy and utility.

HMC is effective to alter mobility profiles constructed from 15 days of mobility. In the next chapter, we search for the vulnerabilities of smaller behaviors (from 1 day down to 30 minutes) and see how current LPPMs including HMC react to them. Towards this purpose, we introduce a novel type of re-identification attacks.

- Chapter 5 -

ILL-Attack: Mobile User Re-identification Using Extremely Randomized Trees

Contents

5.1	Objectives and Roadmap	96
5.2	Problem Illustration	97
5.3	Model of User Re-identification based on Small Traces .	99
5.4	ILL-Attack Design Principles	101
5.4.1	Splitting Algorithm:	101
5.4.2	Data Formatting:	103
5.4.3	ILL-Attack's Classifier: The Extremely Randomized Tree (ERT)	103
5.5	Experimental Evaluation of ILL-Attack	104
5.5.1	Datasets	104
5.5.2	Experimental Setup and Configurations	105
5.5.3	Evaluation of ILL-Attack in a Session-Based Service . . .	106
5.5.4	Evaluation of ILL-Attack in a Crowd-Sensing Application	109
5.6	Conclusion	112

5.1 OBJECTIVES AND ROADMAP

We argue that evaluating the risk of re-identification when sharing data is important for the design of strong privacy preserving mechanisms. In this chapter, we propose ILL-Attack a new re-identification attack that detects the vulnerabilities of re-identification even further than the ones of the state of the art. Indeed, ILL-Attack apprehends differently re-identification. For instance, in the paradigm used in Gambs et al. [30], Primault et al. [82] and Maouche et al. [59] (Chapter 3), mobility data is used to construct user profiles and later on, upon receiving an anonymous trace, the attacker constructs a profile and searches for the most similar one in the past mobility knowledge as a k-NN classifier would do. This type of profiles demand large mobility traces to be applied (i.e., in the order of hours or more), their most common use can be for re-identification in crowd-sensing (traces are in the order of hours or days) or data publishing (traces are in the order of days or weeks). However, for shorter mobility traces (in the order of minutes or few hours), the attacker would try to re-construct the whole profile based on the occurrence of one small behavior. In ILL-Attack, the paradigm is different, the attacker learns from multiple short behaviors in order to be able to recognize them at re-identification. When we talk about "short" or "long" traces, we talk about the time period during which the attacker eavesdropped on the users' mobility or during which the data has been collected by the service provider, we do not talk about the number of records sent by the user or a change in the sampling of the data. As it has been experienced in Section 3.4.9, a profile structure such as the heat map is quite robust to the variations of the proportion of the mobility trace eavesdropped (80% of the maximum user re-identification rate obtained with only eavesdropping randomly 20% records) but for short traces if the heat map does not record a particular behavior, it cannot be used to recognize the behavior later on.

ILL-Attack uses Extremely Randomized Trees to learn users' identity based on their mobility. This attack instantiates a new model of re-identification that divides the mobility traces into multiple shorter mobility traces to learn different behaviors of users in order to be able to re-identify in various scenarios. Its strength is that it can be applied to use cases of a smaller size such as a session-based service (i.e., a user's mobility is collected during the use of a service for a short session), in addition to the

crowd-sensing use case previously studied in the state of the art. In our experiments, we compare ILL-Attack to AP-Attack [59] and POI-Attack [82] in those different use cases on four real mobility datasets. We also study the effectiveness of various LPPMs on ILL-Attack to assess whether current LPPMs can mitigate the threat for user re-identification in the different scenarios.


On average, across all the four datasets, the session-based services (resp. crowd-sensing data) successfully re-identified using mobility data is at 49% (resp. 57%) for ILL-Attack, 34% (resp. 43%) for AP-Attack and 10% (resp. 21%) for POI-Attack. As for the LPPM effects, despite their use, across all the dataset, in the session-based service scenario (resp. crowd-sensing scenario), the rate decrease ranges only from -4% to -14% (resp. from -3% to -18%) for ILL-Attack. And in our experiments, for all the LPPMs, ILL-Attack outperforms AP-Attack and POI-Attack.

Roadmap In the remaining of this chapter, In Section 5.2, we illustrate our motivation, we present the model of re-identification attacks that learn short behaviors and we present the design principles of ILL-Attack. Experimental evaluation results are presented in Section 5.5. And finally, we draw our conclusions in Section 5.6.

5.2 PROBLEM ILLUSTRATION


In this section, we want to showcase the limitations of classical re-identification attacks that uses profiles to re-identify mobility traces in scenarios with short mobility traces (online/semi-online). To do this, we performed an experiment where two state-of-the-art attacks are launched on four mobility datasets in different scenarios depending on the size of the mobility traces. We have a data publishing scenario with full long traces that go up to 5 days (full because for each user in the testing part of the dataset, her whole mobility trace is re-identified as one), a crowd-sensing scenario with medium size traces that go up to 24 hours and a session-based service scenario with short traces smaller than 30min.

The results are shown in Figure 5.1. First, we notice the high results of a data publishing scenario where long mobility traces are available to the attacker (from



(a) Full long traces up to 5 days (e.g., Data Publishing)

(b) Medium size traces up to 24 hours (e.g., Crowdsensing)



(c) Short size traces up to 30 minutes (e.g., Session of an application)

Figure 5.1: Re-identification rates of attacks in different scenarios with different datasets (80% training and 20% testing)



67% up to 74%). The results decrease for a scenario with smaller traces, already with a crowd-sensing scenario where the user sends her data daily (it goes down by -22% on average) and even further with a scenario of short traces (it goes down by an average of -34%).

This could be used as reasoning to claim that smaller mobility traces are not sensitive to re-identification. We argue that this is false and we want to show in this chapter that by changing the way we re-identify mobility traces, an attacker can also break the anonymity of mobility traces in scenarios with short traces. Because, current profile-based attacks aim at reconstructing similar profiles using the mobility traces, while long mobility traces are compatible with this objective, smaller traces have less capacity to reconstruct the profile of the user. To reach this, we use machine learning techniques that learn individual behaviors of users instead of trying to construct a one all-mighty profile to discriminate between users.

5.3 MODEL OF USER RE-IDENTIFICATION BASED ON SMALL TRACES

Consider an attacker that has at his disposal the know user data $\mathbb{KD} = \{T_1, T_2, \dots, T_n\}$. Each T_i is the record of the past mobility of the user U_i . His aim is to use \mathbb{KD} to train a classifier \mathcal{I} that is able upon receiving as input an anonymous mobility trace T to predict the identity of the owner of the mobility trace.

In consequence, the user identities of the system represent the target labels $Y \in \mathbb{U} = \{U_1, U_2, \dots, U_n\}$ and the observation data is $X \in \mathbb{KD}$. X has the form of a time series where each record $r_i = (lat_i, lng_i, t_i)$ is composed of the GPS coordinates with their timestamp. As depicted in Figure 5.2. we first split traces into multiple sub-traces because we need multiple examples for each label (the attacker could receive them split or split them himself to ensure a multitude of examples). This resulting data cannot be fed directly to a classifier learning algorithm. It needs to have the form of a feature vector $X = \{feature1 : v_1, feature2 : v_2, \dots\}$ with $v_k \in \mathbb{R}$. To train the classifier on \mathbb{KD} , we need to use a preparation function \mathcal{T} that transforms $X = \{(lat_1, lng_1, t_1), (lat_2, lng_2, t_2), \dots\}$ to $\mathcal{T}(X) = \{feature1 :$

Figure 5.2: *Re-identification of user by learning on short mobility traces*

$v_1, feature2 : v_2, \dots\}$. This model is generic and different techniques can be applied to achieve such an attack, an instantiation of this model with ILL-Attack is presented in the next section. To summarize in order to design a re-identification attack as a multi-label classifier, we need :

- A splitting algorithm that turns a mobility trace into multiple sub-traces.

$$\begin{aligned} \mathcal{S} : (\mathbb{R}^2 \times \mathbb{R}^+)^* &\rightarrow ((\mathbb{R}^2 \times \mathbb{R}^+)^*)^* \\ T = \{r_1, r_2, \dots\} &\mapsto \mathcal{S}(T) = \{\{r_1, \dots, r_{k_1}\}, \{r_{k_1+1}, \dots, r_{k_2}\}, \dots\} \end{aligned} \quad (5.1)$$

- A transformation function that turns a mobility sub-trace into a feature vector of dimension d .

$$\begin{aligned} \mathcal{T} : (\mathbb{R}^2 \times \mathbb{R}^+)^* &\rightarrow \mathbb{R}^d \\ T = \{r_1, r_2, \dots\} &\mapsto \mathcal{T}(T) = \{v_1, v_2, \dots, v_d\} \end{aligned} \quad (5.2)$$

- A multi-label classifier.

$$\begin{aligned} \mathcal{I} : (\mathbb{R}^d \times \mathbb{U})^* \times \mathbb{R}^d &\rightarrow \mathbb{U} \\ (\mathcal{T}(\mathbb{KD}), \mathcal{ID}(\mathbb{KD}), \mathcal{T}(T)) &\mapsto \mathcal{I}(\mathcal{T}(\mathbb{KD}), \mathcal{ID}(\mathbb{KD}), \mathcal{T}(T)) = U_a \end{aligned} \quad (5.3)$$

5.4 ILL-ATTACK DESIGN PRINCIPLES

We use the previous model to design ILL-Attack (Identity Learning with Location Attack) an attack based on fixed slices of sub-traces and Extremely Randomized Trees (ERT) [32].

5.4.1 Splitting Algorithm:

this algorithm takes fixed time slices of length Δ (in seconds) from the mobility trace. As described in Algorithm 4, each time interval $[t_0 + k\Delta, t_0 + (k+1)\Delta[$ corresponds to one sub-trace. The value of Δ should be chosen to ensure multiple examples for the training phase. It should also correspond to the size of the mobility trace to re-identify.

Algorithm 4 Splitting Algorithm of Long Mobility Traces in ILL-Attack

```

1: function  $\mathcal{S}_f(T, \Delta)$ 
2:    $slices \leftarrow \emptyset$  Set of slices to output
3:    $currentSlice \leftarrow \emptyset$ 
4:    $st \leftarrow T[0]^{(t)}$  Current slice starting time
5:   for  $r$  in  $T$  do Consider each record in the trace
6:     if  $r^{(t)} < (st + \Delta)$  then Current slice cannot exceed  $\Delta$ 
7:       Continue to gather records for this slice
8:        $currentSlice \leftarrow currentSlice \cup \{r\}$ 
9:     else
10:       $slices \leftarrow slices \cup currentSlice$  Save the slice
11:       $currentSlice \leftarrow \{r\}$  Start a new slice
12:       $st \leftarrow st + \delta$  Update the starting time of the new current slice
13:    end if
14:  end for
15:   $slices \leftarrow slices \cup currentSlice$ 
16:  return  $slices$ 
17: end function

```

Feature	Description	Value Range
<i>hourOfDay</i>	Average hour of the day of recording	$[0, 23]$
<i>nbRecords</i>	The number of records in this sub-trace	\mathbb{R}^+
<i>centerLat</i>	The latitude of the centroid of the trace	$[-90, 90]$
<i>centerLng</i>	The longitude of the centroid of the trace	$[-180, 180]$
<i>IDofCell_i</i>	The Proportion of records in cell i for all non-empty cells in the dataset	$[0, 1]$

Table 5.1: *List of features used by ILL-Attack*

5.4.2 Data Formatting:

ILL-Attack mainly uses the heat map as a way of representing the whole mobility trace with a fixed dimension space. The map is divided into regions of the same size. In each cell, we compute the proportion of records in it.

In addition to the heat map, other features are added to each observation as described in Table 5.1. Specifically, we use *hourOfDay*, which adds temporal information on the trace to differentiate similar moving patterns between day-night shifts, e.g., a user living near the working place of another user. They may have similar heat map prints but at different times of the day (this feature is limited to sub-traces that are smaller than a day). We also use the *nbRecords* feature, which gives information about the sub-trace used for the construction of the heat map. Finally, we use the centroid of the sub-trace, which gives the average position of the user on the map.

5.4.3 ILL-Attack's Classifier: The Extremely Randomized Tree (ERT)

ILL-Attack uses to train its model an Extremely Randomized Tree classifier [32]. It constructs a set of M decision trees where the splitting feature chosen is the best one among K randomly selected attributes and the selected cutting point for each feature is random. The main advantage of this type of classifier compared to other famous decision tree classifiers such as Random Forest is the fact that it eliminates

the burden of searching for the optimal cut-point without deteriorating the accuracy of the classifier. ERT was the best method found using an evolutionary algorithm searching for the most suitable classifier for our task (with cross-validation to avoid over-fitting) [75].

5.5 EXPERIMENTAL EVALUATION OF ILL-ATTACK

In the following, we first present the real-life mobility datasets used in our experiments (Section 5.5.1). Then, we describe the experimental environment and configuration settings we used (Section 5.5.2). We compare ILL-Attack to state-of-the-art attacks in two use cases: (1) Re-identifying a short session of a service using mobility data in Section 5.5.3. (2) A crowd-sensing campaign where the users send their data daily in Section 5.5.4. The attacks are evaluated in both a non-obfuscated (NOBF for short) setup and an obfuscated setup using state-of-the-art Location Privacy Protection Mechanisms (LPPMs).

5.5.1 Datasets

We used four real mobility datasets in our experiments. These datasets are: (1) Cabspotting [80] that contains the mobility of 536 cab drivers in the city of San Francisco; (2) Geolife [106] that contains the mobility of 42 users mainly in the city of Beijing; (3) MDC [54] that contains the mobility data of 144 users in the city of Geneva and (4) PrivaMov [9] that contains the mobility of 48 students and staff members in the city of Lyon. A mobility trace is constituted of a sequence of spatio-temporal records $r = (lat, lng, t)$ associated to a given user, where lat and lng correspond to the latitude and longitude of GPS coordinates while t is a timestamp.

To make the comparison fair between the various datasets, we selected in each dataset the 30 most active successive days. We also took a subset of the 50 most active users of cabspotting to lower the memory usage. We present in Table 5.2 a description of the datasets used in our experiments. The users are not active in all the days of the period, some are more active than others. We consider as a mobility trace, the mobility of the user during all the period.

Table 5.2: *Description of the datasets used for the experiments on ILL-Attack*

Dataset	CabSpotting	Geolife	MDC	PrivaMov
#users	536 (50)	42	144	48
Location	San Francisco	Beijing	Geneva	Lyon
#records	11 219 955 (1 409 687)	1 574 338	904 422	973 684

5.5.2 Experimental Setup and Configurations

The following experiments were conducted in a computer running an Ubuntu 14.04 OS with 50GB of RAM and 16 cores of 1.2Ghz each. ILL-Attack's prototype is developed in Python and is available for download at: <https://github.com/mmaouche-insa/ILL-Attack>. It uses scikit-learn [78] for the Extremely Randomized Tree implementation and S2-Geometry library [88] for the decomposition of the map into areas of equal size.

ILL-Attack parameter is the size of the cell in the S2-Geometry library, which is set at level 13 meaning that an area covers approximately a block of a city. For the Extremely randomized tree, the number of decision trees is set to 100, the criteria to measure the quality of a split is the Gini criteria [10]. The minimum number of samples required to be at a leaf node is set to 2, the minimum number of samples required to split an internal node is set to 8. Finally, the number of features to consider when looking for the best split is 60%. This configuration was searched using an automated tool Tpot [75] that uses evolutionary algorithms to search the parameters' space. It has been launched only once on the GeoLife dataset and not reconfigured on each dataset or use case. POI-Attack has two parameters for the extraction of POIs from mobility traces. These parameters are the diameter of the clustering area, and the minimum time spent inside a POI. These parameters are respectively set to 500 meters and 15 minutes to accommodate small traces (for the first use case for example). AP-Attack has a configuration parameter that corresponds to the square cell size, which was set to 800 meters [59]. The LPPMs come with their own configuration parameters, that are set as follows. Geo-I's ϵ configuration parameter is set to 0.01; this adds a medium amount of noise to the obfuscated data

(the lower ϵ the higher the noise). Promesse's α configuration parameter is set to 200 meters, it represents the distance between two successive sampling points (default value). HMC's main parameter is the square cell sizes of the heat map and it was set to 800 meters (default value). Finally, TriLateralation's parameter is the radius of the circular region where random points are generated, which was set to $1km$ (default value).

5.5.3 Evaluation of ILL-Attack in a Session-Based Service

In this section, we present a description of the session-based service, followed by the evaluation of ILL-Attack compared to two state-of-the-art re-identification attacks then we show the resilience of four LPPMs against ILL-Attack.

Description:

In this use case, we consider a session re-identification of a location service. Let us assume a user that uses an application for a short period of time (e.g., short navigation in the city, successive location check-ins, multiple searches for restaurants. . .). During this usage, the user sends her location multiple times during a single session but between two sessions the user changes her ID in the system (e.g., using a privacy proxy or an anonymity network such as TOR). As a consequence, the attacker can know that this sequence of records came from the same user but the attacker cannot reassemble all the small traces of the user as one big trace since each of them has a different id.

In this use case, the attacker has access to all the 80% past mobility of the users (i.e., \mathbb{KD}). AP-Attack and POI-Attack use this mobility traces to construct the profiles, while ILL-Attack uses the splitting algorithm with $\Delta = 30mn$ to train its classifier. In this use case, for all the attacks, the attacker receives a mobility trace of one session usage, which is simulated in our experiment by slices of sub-traces smaller than 30 minutes (hence the Δ for ILL-Attack).




Figure 5.3: *ILL-Attack compared to two state-of-the-art attacks in the session-based service scenario (30mn sessions – 80% training and 20% testing)*

Performance of ILL-Attack in Session-based Services:

In Figure 5.3, we compare ILL-Attack against two re-identification attacks AP-Attack and POI-Attack. The number of re-identifications attempted by the attacker (i.e., size of the test set) is respectively 2437 for Cabspotting, 941 Geolife, 33892 for MDC and 1571 for Privamov. The results show that ILL-Attack outperforms state-of-the-art attacks. On average across all the sessions of all the dataset, ILL-Attack successfully re-identified 49% of the sessions, while AP-Attack re-identified 34% and POI-Attack 10%. It is important to mention that these sessions are quite short in terms of duration. A small sub-trace of a trajectory shared by a lot of users in the system is hard to re-identify. The attacks based on profile construction such as AP-Attack and POI-Attack try to search for a corresponding profile by re-constructing the profile using a short sub-trace. This is the reason why AP-Attack has lower performance than ILL-Attack. Because rather than trying to re-construct a whole profile, ILL-Attack learns from individual short past behaviors and for re-identification, it searches for the identity associated with such behavior. While POI-Attack has even worse performance than AP-Attack since the latter better captures the profiles as it has been shown in the previous chapters.

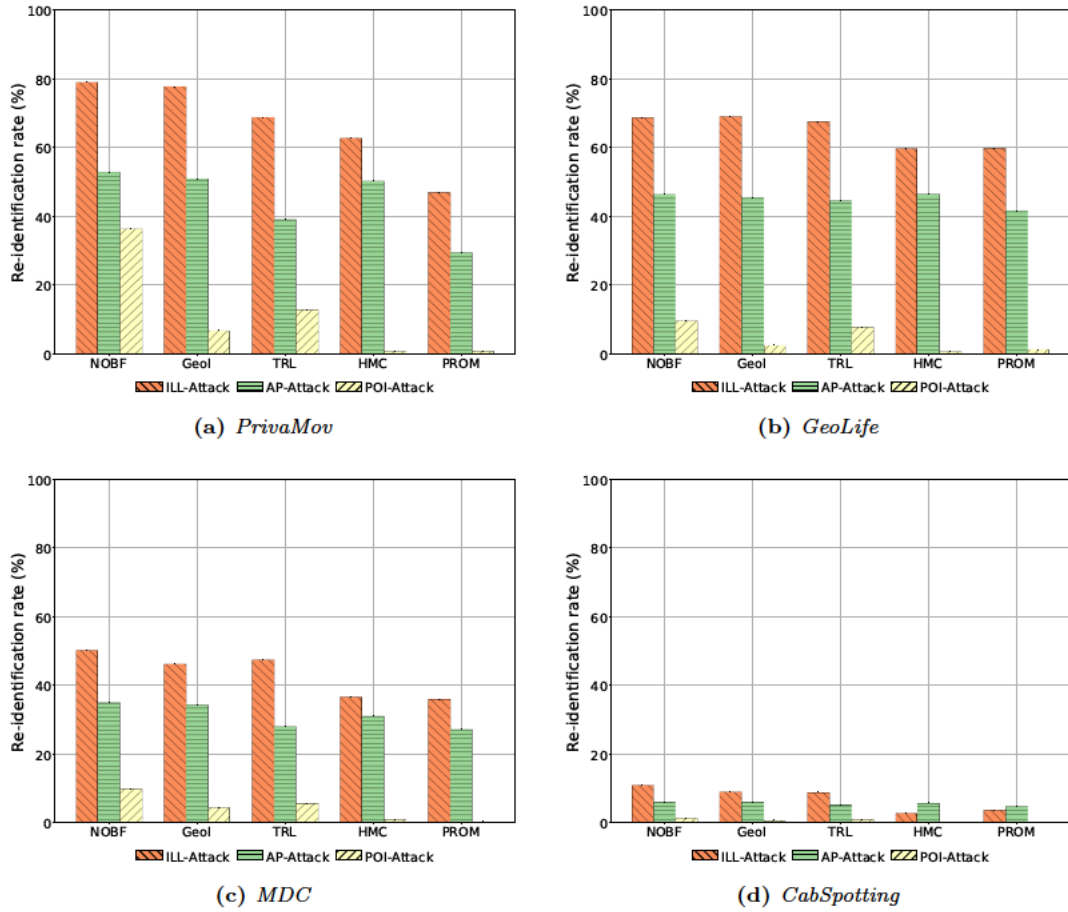


Figure 5.4: Resilience of ILL-Attack against LPPMs compared to two state-of-the-art attacks in the session-based service scenario (30mn sessions – 80% training and 20% testing)

LPPMs Effectiveness Against Re-identification Attacks in Session-based Services:

In this experiment, we consider that each session the attacker wants to re-identify has been previously obfuscated by an LPPM. If the LPPM is an online LPPM the records are obfuscated one by one and if the LPPM is semi-online or offline the whole mobility trace of the session is obfuscated before being received by the attacker (the training dataset is never obfuscated by any LPPM).

In Figure 5.4, we compare the LPPM effectiveness against re-identification attacks. We first notice that the LPPM have low effectiveness on these short sub-traces ($< 30min$). Promesse and Trilateration seem to be the LPPMs with the most effect even though they do not protect the sessions that much. Since, on ILL-Attack, even the most effective LPPM decreases the user re-identification rate by only -14% . If we compare ILL-Attack, AP-Attack and POI-Attack after the usage of LPPM, ILL-Attack outperforms the other attack in all the datasets and against all the LPPMs. If we consider all the users of all the datasets, ILL-Attack's re-identification rates are still between 35% and 46%, while AP-Attack's rates are between 26% and 33% and POI-Attack's rates are between 0% and 8%.

5.5.4 Evaluation of ILL-Attack in a Crowd-Sensing Application

In this section, we present a description of this use case, followed by the evaluation of ILL-Attack compared to AP-Attack and POI-Attack then we show the resilience of ILL-Attack against GeoI, TriLateration, HMC and Promesse.

Description:

In this use case, we consider a crowd-sensing scenario. The user gathers data during the day, then at night when her mobile phone is connected to a WIFI connection and put to charge, she sends all the data gathered during the day to either the crowd-sensing platform or a privacy proxy if need be. In this use case, the attacker has access to at most a mobility trace of a whole day but cannot reassemble all the




Figure 5.5: *ILL-Attack compared to two state-of-the-art attacks in the crowd-sensing scenario (24h data – 80% training and 20% testing)*

mobility traces gathered by the same user during the crowd-sensing campaign since each data portion sent has a different ID.

In this use case, the attacker has access to all the 80% past mobility of the users (i.e., \mathbb{KD}). AP-Attack and POI-Attack use this mobility traces to construct the profiles while ILL-Attack uses the splitting algorithm with a Δ of 24 hours to train its classifier. In this use case, for all the attacks, the attacker tries to re-identify one portion of data sent by the user, which is simulated in our experiment by slices of sub-traces smaller than a day (hence the Δ for ILL-Attack).

Performance of ILL-Attack in a Crowd-sensing Application:

In Figure 5.5, we compare ILL-Attack against two re-identification attacks AP-Attack and POI-Attack. The number of re-identifications attempted by the attacker (i.e., size of the test set) is respectively of 200 for Cabspotting, 85 Geolife, 432 for MDC and 82 for Privamov. The results show that ILL-Attack outperforms state-of-the-art attacks. On average across all the sessions of all the dataset, for ILL-Attack 57% of the portion of data are re-identified, while AP-Attack re-identifies 43% and POI-Attack 21%. We first notice that the success rates are higher than the use case of session usage. Even though the amount of data the attackers learn from and tries to re-identify on is the same (80% train and 20% test), the test is sub-divided in a smaller number of sub-traces that have a bigger number of records ($< 24h$ against

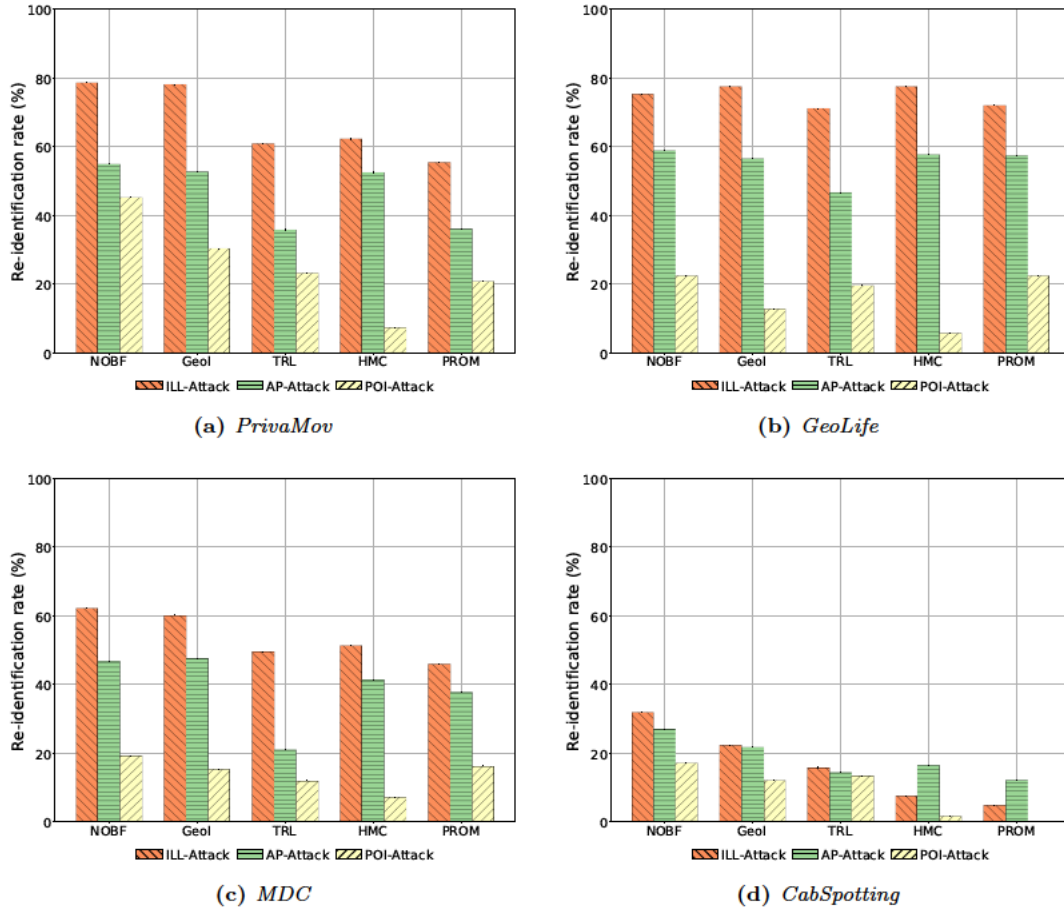


Figure 5.6: Resilience of ILL-Attack against LPPMs compared to two state-of-the-art attacks in the crowd-sensing scenario (24h data – 80% training and 20% testing)

$< 30min$). This is particularly true for Cabspotting (11% to 31%), since the small sub-traces under $30min$ have a higher chance of corresponding to a frequent route of a cab driver shared by others, while a 24h-mobility captures more the private discriminative patterns of a cab driver.

LPPMs Effectiveness Against Re-identification Attacks in a Crowd-sensing Application:

In Figure 5.4, we compare the LPPMs' effectiveness against re-identification attacks. We first notice that on average the LPPMs have a higher effectiveness than in the

first use case. This is especially the case for HMC that goes up from a decreasing effect of -3% in the first use case to -13% for ILL-Attack across all the sub-traces of all the datasets. Against ILL-Attack, Promesse has the higher effect with a decrease of -18% of all the sub-traces across the dataset re-identified. Trilateration and HMC have -13% . All in all, ILL-Attack outperforms AP-Attack and POI-Attack even after the protection of the LPPMs.

5.6 CONCLUSION

In this chapter, we presented ILL-Attack a novel re-identification attack that uses machine learning to learn the short behaviors of users. This attack uses the heat map representation of a mobility trace and the timestamp of the mobility trace to differentiate user patterns. ILL-Attack was evaluated on four real mobility datasets against three representative LPPMs and compared to two classical state-of-the-art re-identification attacks that construct user profiles. The evaluation was conducted on two use cases, one where an attacker aims at re-identifying session-based services usage through mobility data and the other where an attacker re-identifies crowd-sensing data sent by a user on a daily basis. The results show that ILL-Attack outperforms the other two attacks in all the use cases and against all the LPPMs proposed.

As future work, we consider using machine learning techniques to protect users against re-identification attacks. Either by filtering discriminative behavior or generating traces that preserve utility while hiding sensitive and discriminative moving patterns. A good perspective to investigate is the utilization of Generative Adversarial Networks that might be able to respond to the threat of re-identification.

- Chapter 6 -

Hybrid-LPPM: A User-Centric Fine-Grained Multi-LPPM

Contents

6.1	Objectives and Roadmap	114
6.2	Problem statement	115
6.3	Design of <i>Hybrid-LPPM</i>	117
6.4	Experimental Evaluation of <i>Hybrid-LPPM</i>	118
6.4.1	Experimental Setup, Configurations and Datasets	118
6.4.2	Privacy and Utility Evaluation	119
6.5	Conclusion	121

6.1 OBJECTIVES AND ROADMAP

According to the previous results of Chapter 3 and Chapter 5, we notice that the current LPPMs applied to user mobility data (whether in the form of session-based services, in the form of longer crowd-sensing data portions or even in a data publishing scenario) are not sufficient to protect users in front of user re-identification attacks.

Indeed, after analyzing the previous results more thoroughly, we noticed that users are affected differently by the LPPMs and that even portions of mobility data of the same user are not equally protected by the same LPPM. This shows the need to develop new LPPMs that are able to face the threat previously illustrated in all the evaluations (ILL-Attack and AP-Attack more particularly). As a first attempt, we propose to consider the particularity of each behavior of the user and design LPPMs that change their obfuscation depending on the sub-trace. This is why we propose to make use of off-the-shelf state-of-the-art LPPMs and apply the best one for each sub-trace.

We propose *Hybrid-LPPM* that operates in a crowd-sensing application where the user goes through a privacy proxy each time it needs to send a mobility trace to the analyst. The privacy proxy is first responsible for hiding the ID of the user and hiding the source of the data. Also, it uses background knowledge sent by different users to choose the best LPPM to apply for this particular sub-trace. To avoid information leakage, the proxy could rely on secure hardware such as trusted execution environments. This technology is being increasingly available on commodity hardware (e.g., the sky lake generation of intel processors with SGX [92]). It also operates in a data publishing scenario, the publisher should then use *Hybrid-LPPM* instead of experimenting on each LPPM individually.

The best LPPM is chosen using two criteria: (1) Privacy: we choose the set of LPPMs that protect the most against a set of re-identification attacks trained with the gathered background knowledge. (2) Utility: from those LPPMs, we select the one with the best utility according to a chosen metric (we evaluate how much the data has been distorted).

Roadmap During the remaining of this chapter, we start by illustrating the problem in Section 6.2. We describe our method in Section 6.3. We present the evaluation in Section 6.4. We conclude in Section 6.5.

6.2 PROBLEM STATEMENT

A user mobility trace can be obfuscated using different LPPMs available in the state-of-the-art. Each one has different results on different users according to the attack. For example, in Table 6.1, we show the re-identification results of a sample of 24 hours long sub-traces of the dataset PrivaMov obfuscated with four different LPPMs. We notice that there are sub-traces where only a specific LPPM is able to protect the user against all the attacks (HMC for trace 42-23 and Promesse for 42-24 and 27-13). Moreover, this specific LPPM varies between users and even between sub-traces of the same user (user 42 needs HMC for her sub-trace 23, while she needs Promesse for her sub-trace 24). In other cases, multiple LPPMs can be chosen (the sub-trace 75-25 can use GeoI, Promesse or HMC and 50-28 can use any of the four LPPMs). But since each LPPM has different utility values, it would be wise to choose the one that deteriorates the utility the less.

These are only examples, we notice this phenomenon across all datasets as shown in Figure 6.1. We describe for each dataset the protection results when the sub-traces are obfuscated with the trio Promesse, TriLateration and HMC. We consider only the sub-traces protected against the three re-identification ILL-Attack, AP-Attack and POI-Attack simultaneously. We separate those protected sub-traces first according to the number of LPPMs that protected them, then according to the set (e.g., from the trio three different pairs could be the one protecting the sub-trace). The results show that there is no one-size-fits-all LPPM and that depending on the sub-trace a different set of LPPM should be considered. To make those choices we use a privacy proxy and we describe how it operates in the next section.

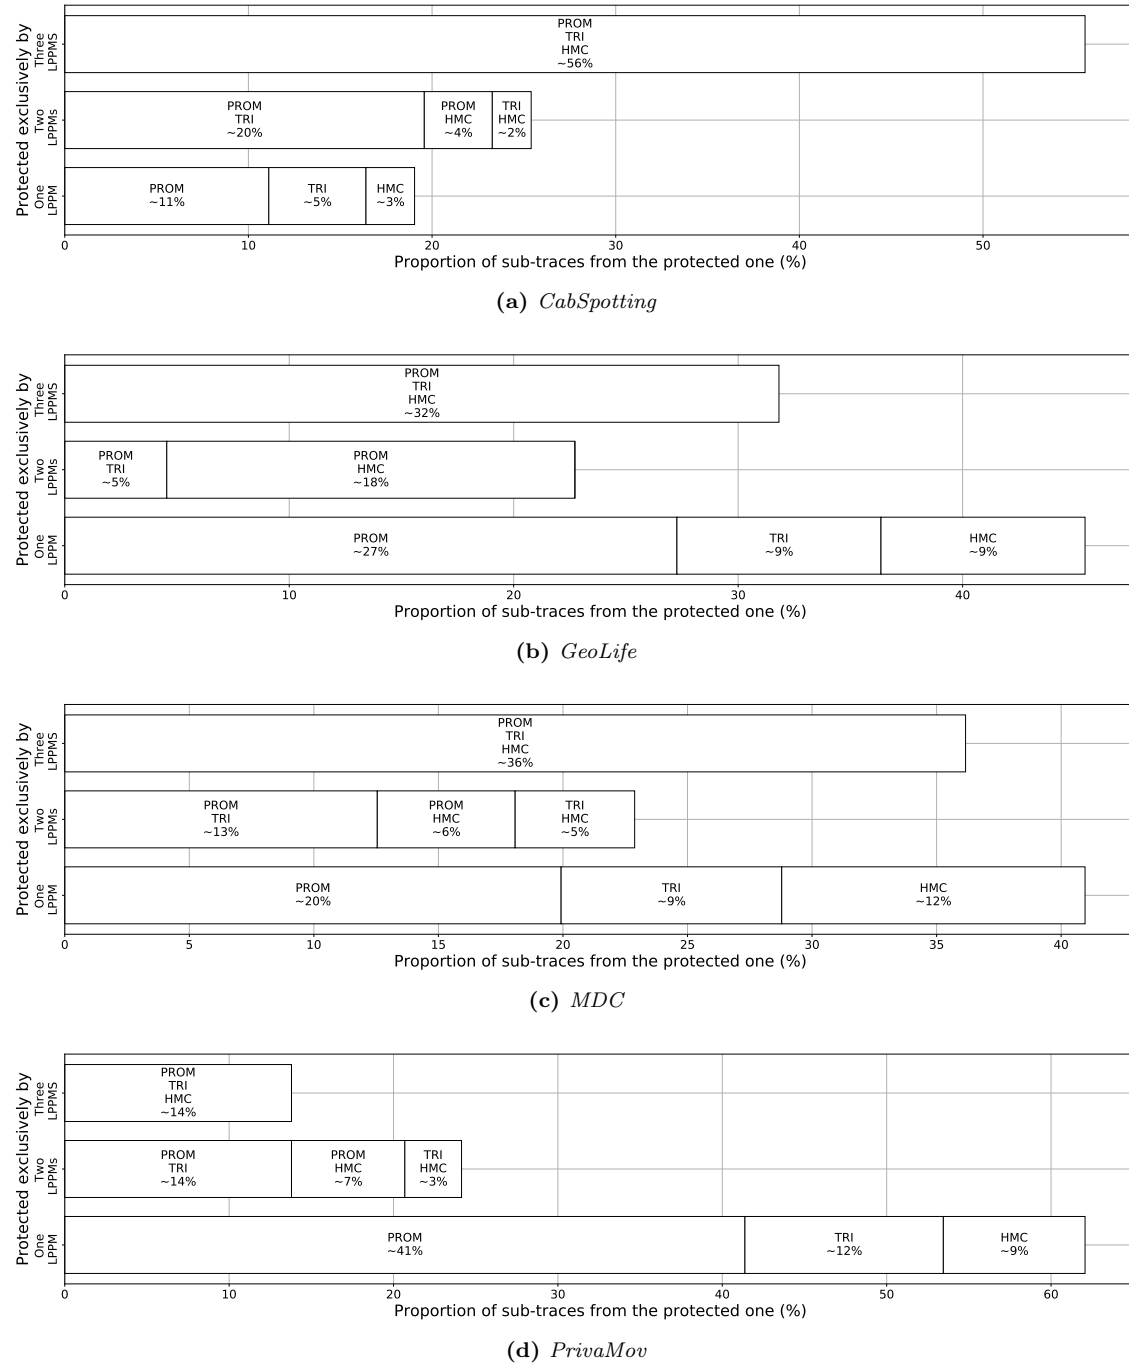


Figure 6.1: *Illustration of how sub-traces have different subsets of LPPMs that protects them*

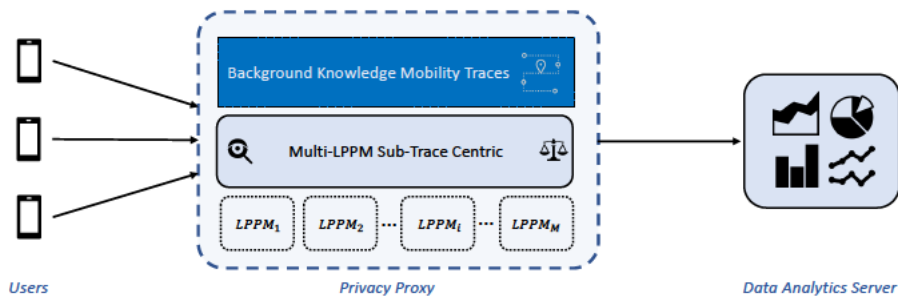
Table 6.1: Illustration of how different sub-traces have different re-identification results using a sample of sub-traces of the *PrivaMov Dataset*

User ID	Sub-trace N°	GeoI			Promesse			TriLateration			HMC		
		ILL	AP	POI	ILL	AP	POI	ILL	AP	POI	ILL	AP	POI
42	23	X			X			X					
42	24	X						X			X		
27	13	X	X	X				X	X		X	X	X
75	25							X					
50	28												

X: User correctly Re-identified - The LPPM Fails

6.3 DESIGN OF *Hybrid-LPPM*

In Figure 6.2, we present the architecture of the privacy proxy. It uses background knowledge of all the users in order to evaluate the risk of re-identification and have multiple LPPMs to use in order to protect the user mobility traces. A user sends her mobility trace and the privacy proxy chooses the most suitable LPPM to apply and sends its results to the data analytics server. This also applies in a data publishing scenario, the publishers apply the hybrid solution instead of experimenting on multiple LPPMs. To choose the most suitable LPPM to apply from a list of available LPPMs $\mathbb{L} = \{\mathcal{L}_1, \mathcal{L}_2, \dots, \mathcal{L}_K\}$. *Hybrid-LPPM* uses multiple re-identification attacks $\mathbb{A} = \{\mathcal{A}_1, \mathcal{A}_2, \dots, \mathcal{A}_K\}$ and their weights $\mathbb{W}^A = \{w_1^A, w_2^A, \dots, w_K^A\}$. The Privacy risk offered by an LPPM \mathcal{L} on a mobility trace T would be evaluated using the formula

**Figure 6.2:** Fine-grained protection mechanism with the sub-trace centric approach of LPPMs' protection

of Equation 6.1. Higher the value, higher the risk of a user being re-identified by an attacker.

$$\mathcal{P}_{\mathcal{L}}(T) = \frac{\sum_{\mathcal{A}_i} \begin{cases} w_i^{\mathcal{A}} & \text{If } \mathcal{A}_i(\mathcal{L}(T), \mathbb{K}\mathbb{D}) = \mathcal{ID}(T) \\ 0 & \text{Else} \end{cases}}{\sum_{\mathcal{A}_j} w_j^{\mathcal{A}}} \quad (6.1)$$

The priority is given to privacy. So, from the set of available LPPMs \mathbb{L} a sub-set $\mathbb{L}_{min}(T)$ with minimum privacy risk for T is chosen (See Eq.6.2)

$$\mathbb{L}_{min}(T) = \left\{ \mathcal{L} \in \mathbb{L} \mid \mathcal{P}_{\mathcal{L}}(T) = \min_{\mathcal{L}_i} \mathcal{P}_{\mathcal{L}_i}(T) \right\} \quad (6.2)$$

To choose the most suitable LPPM to apply, we use a utility metric \mathcal{UT} . The most suitable LPPM \mathcal{L}^* is the LPPM from $\mathbb{L}_{min}(T)$ with the best utility (as described in Equation 6.3). The utility evaluation can be defined using multiple metrics similar to how \mathcal{P} was defined in Eq.6.1.

$$\mathcal{L}^* = \arg \max_{\mathcal{L} \in \mathbb{L}_{min}(T)} \mathcal{UT}(T, \mathcal{L}(T)) \quad (6.3)$$

6.4 EXPERIMENTAL EVALUATION OF *Hybrid-LPPM*

6.4.1 Experimental Setup, Configurations and Datasets

The following experiments were conducted in a computer running an Ubuntu 14.04 OS with 50GB of RAM and 16 cores of 1.2Ghz each. The re-identification attacks and the LPPMs are configured in the same way as the one of the experiments of Chapter 5 (see Section 5.5.2 for more details) To summarize, we use three re-identification attacks for the adversary. ILL-Attack with regions of level 13, AP-Attack with square cells of size 800 meters and POI-Attack with a clustering algorithm parameterized with 500 meters for the maximal clustering area and 15 minutes for the minimal duration. For *Hybrid-LPPM*, we give to each attack a weight of 1, which makes the privacy evaluation of the proxy similar to the number of successful attacks metric used for

the experiments of HMC in Chapter 4. We use one utility metric, the Area Coverage with square cells of size 800 meters. We use four LPPMs, Geo-I with $\epsilon = 0.01$, Promesse with $\alpha = 200$ meters, TriLateralation with a radius of $1km$ and HMC with square cells of size 800 meters. We used the same four real mobility datasets in our experiments as the one for the evaluation of ILL-Attack (Section 5.5.1).

6.4.2 Privacy and Utility Evaluation

We evaluate the hybrid scenario in a crowd-sensing setting similar to the one of the evaluation of ILL-Attack in the previous chapter. We consider a crowd-sensing application where users send their data daily (every 24 hours). Each data is sent to the privacy proxy that redirects them to the crowd-sensing platform. The privacy proxy applies the hybrid method using a background knowledge shared by all the users in a secure privacy proxy. We consider a strong attacker that has access to the totality of the background knowledge of the privacy proxy rather than small leakages from the previous mobility of users. We can assume that the attacker or too curious analyst used to receive mobility data of the users.

The results are depicted in Figure 6.3. As expected *Hybrid-LPPM* strictly ameliorates all the results of the individual LPPMs. If we consider all the 799 sub-traces across all the datasets, 68% of the sub-traces are fully protected, while for the single LPPMs it ranges from 35% to 54%. Also, *Hybrid-LPPM* increases the proportion of sub-traces fully protected with an Area Coverage above 0.75 to 48%, while the proportion ranges from 1% to 35% for the other LPPMs.

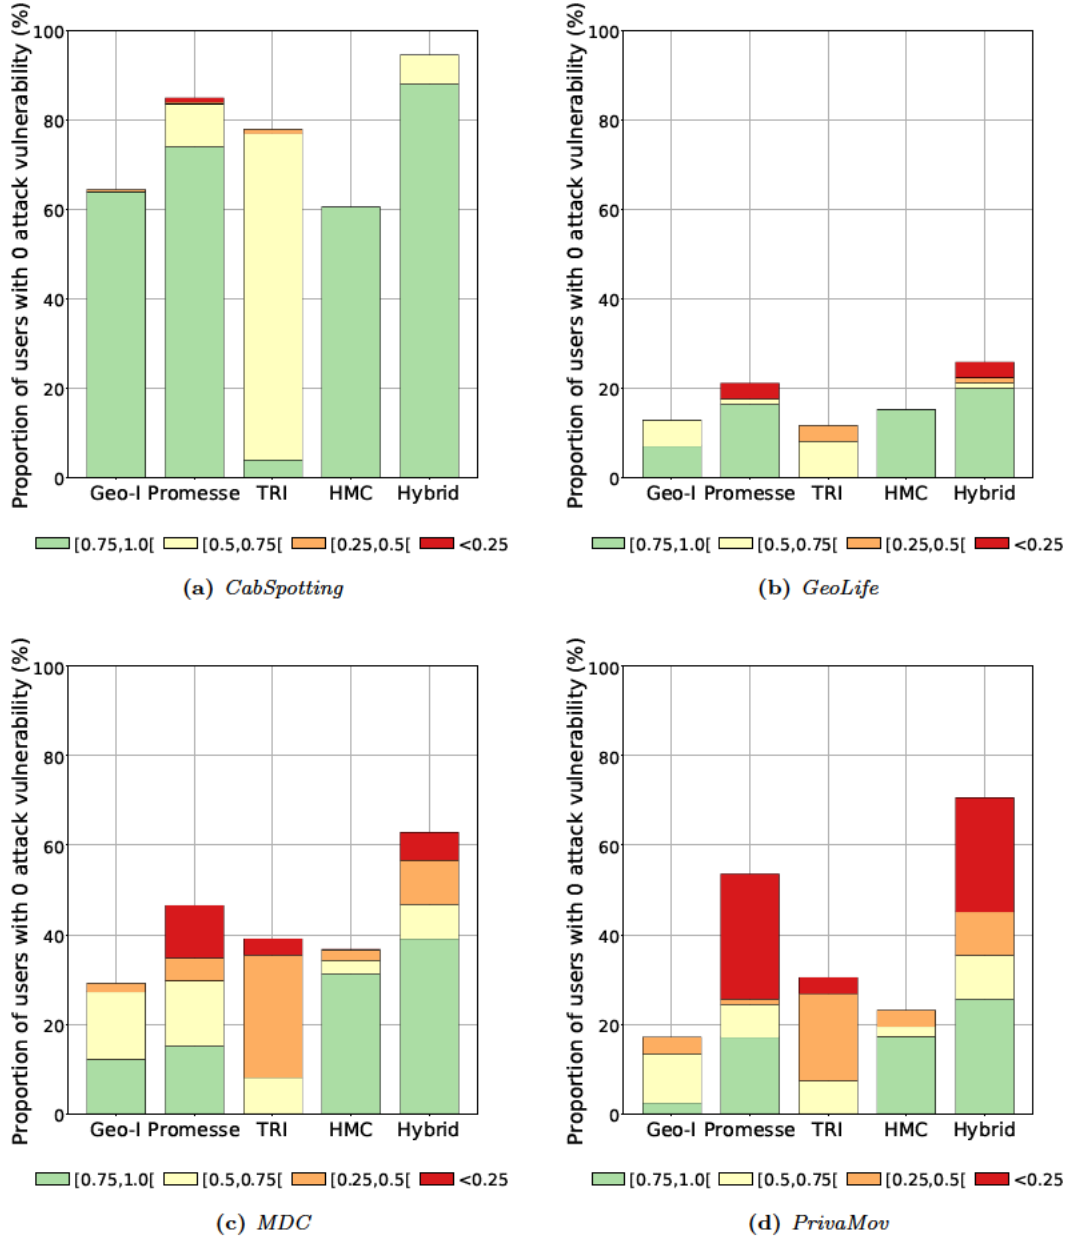


Figure 6.3: Comparison of Hybrid-LPPM with single LPPM competitors - Area coverage utility metric - 80 % for training and 20 % for testing

6.5 CONCLUSION

In this chapter, we propose a solution to make use of off-the-shelf LPPMs to mitigate the risk of re-identification in the context of a crowd-sensing application or a data publishing scenario. Our method *Hybrid-LPPM* first gives priority to privacy by minimizing the number of successful re-identification attacks than chooses the LPPM with the best utility among the most protective ones. We experimented *Hybrid-LPPM* on four real mobility datasets against three re-identification attacks and compared it to the protection of four individual LPPMs. Across all the 799 sub-traces in the testing set, 68% of them are fully protected against all the three re-identification attacks (bigger than its competitors with 35% up to 54%).

- Chapter 7 -

Conclusion & Perspectives

7.1 CONCLUDING REMARKS

With the wide usage of location-based services, a great amount of mobility data has been received by different service providers. Even though awareness on privacy preservation is raising, anonymizing mobility data from now on is ineffective since mobility data is unique and thus discriminative. An adversary that has access to past mobility of users can re-associate anonymous mobility traces to known mobility using re-identification attacks.

In this thesis, we tackled the threat of re-identification attacks on mobility data. Our goal was to discover vulnerabilities of current systems and to design privacy mechanisms able to counter this threat. We took particular care of differentiating long traces with multiple occurrences of discriminative patterns and short traces that may seem hard to re-identify that are actually also at risk. When protecting mobility traces, we consider traces as macro-mobility and we aim at modifying the intrinsic behavior of users, rather than focusing on particular discriminative aspects such as POIs. In the contributions of this thesis, the utility was put on a pedestal. Our protection mechanisms were guided by various utility metrics for various applications.

In this thesis, we started by **constructing a re-identification attack named**

AP-Attack with robust profiling of users based on heat maps, a spatial aggregation of a user mobility trace in square regions of the map. We also propose a novel paradigm of re-identification attack where the attacker does not consider only one possible identity as an output of a re-identification attack but rather considers multiple identities depending on a **re-identification policy**. The goal of a policy is to have a selection of a small number of identities that need to be further investigated. The attacker aims at having the smallest set of possible identities while including the correct identity. In consequence, we propose new ways to measure the strength of an attacker by considering the set size of possible identities and the number of false positives.

In order to propose a countermeasure against profile-based re-identification attacks, **HMC a utility constrained LPPM for crowd-sourcing and data publishing** was proposed. This LPPM reasons on user mobility as a whole, captured using heat maps. HMC extracts user profiles by constructing heat maps and alters it by making it look similar to the heat map of another user. To limit the decrease in data utility, HMC uses the heat map of the closest user as a basis for performing the alteration. and it transforms back each altered heat map to a mobility trace by trying to retain as much as possible the users' original traces unchanged.

We also tackle the re-identification of short traces that may seem protected since an attacker would build profiles using them with difficulty. But we show the risk of **re-identification of short traces using multi-trace learning rather than profiling with ILL-Attack**. It learns from multiple short behaviors in order to be able to recognize them at re-identification. ILL-Attack uses Extremely Randomized Trees to learn users' identity based on their mobility. This attack instantiates a new model of re-identification that divides the mobility traces into multiple shorter mobility traces to learn different behaviors of users in order to be able to re-identify in various scenarios. Its strength is that it can be applied to use cases of smaller size.

In this thesis, we propose a method to make use of off-the-shelf LPPMs with **Hybrid-LPPM a user-centric fine-grained multi-LPPM**. Because users are affected differently by the LPPMs and that even portions of mobility data of the same user are not equally protected by the same LPPM. *Hybrid-LPPM* operates in a crowd-sensing application where the user goes through a privacy proxy each time it needs to send a mobility trace. The privacy proxy is first responsible for

hiding the ID of the user and hiding the source of the data. Also, it uses background knowledge sent by different users to choose the best LPPM to apply for this particular sub-trace. The best LPPM is first chosen as the one that protects the most against a set of re-identification attacks trained with the gathered background knowledge and secondly it is selected as the one with the best utility according to a chosen metric (i.e., privacy then utility).

7.2 FUTURE WORK

Enforcing Formal Guarantees with HMC

HMC protects users against re-identification attacks by transforming heat maps (i.e., a probability distribution) to make them look similar to the one of other users. Two approaches are possible extensions:

k-anonymity: currently HMC obfuscates data of one user using the past mobility of another user. As a direct extension, it could use current data of one user and provide 2-anonymity. It could also consider the heat map of multiple users and build a cluster of k users to enforce k -anonymity. We would construct a centroid profile of the k users inside the cluster and transform each profile to look similar to the average profile of the cluster.

Differential Privacy: heat maps are used for re-identification since they permit to enclose POIs, paths and absences of POIs in one data structure. With a divergence such as the Topsoe diverge to measure the dissimilarity between them. A possible improvement for the protection of HMC is to make the obfuscated heat map differentially private so as to not disclose too much information to an attacker. Using an extension of Geo-Indistinguishably, rather than using it on geo-located records, we would use on our high dimensional heat map. To bound the information an attacker might get for the heat map especially the closeness to the user past mobility heat map.

Generative Adversarial Networks (GANs) to protect users against re-identification attacks

GANs [35] are frameworks in which two neural networks compete, the first network called *Generator* (G) generates data that captures the training data distribution and another network called *Discriminator* (D) that estimates the probability that a sample came from the training data rather than G . In the ongoing work of Romanelli et al. [87], they took inspiration of GANs to propose a method to generate noise for location data using two networks a Generator that generates the noise and a Classifier that aims at distinguishing between the different users (they do not consider re-identification with past mobility as in this thesis). A possible future work for this thesis would be to also take inspiration with GANs in order to protect mobility data against re-identification attacks (with past mobility).

Creating Fake Profiles

Currently, HMC needs to gather user past data. As future work, the extension of using fake profiles as the target users for the confusion would be interesting. First, since user may not accept that their data is used for the protection of other users. Secondly, in order to make HMC independent of the mobility of other users and independent of a privacy proxy. We could leverage synthetic traces of fake users to be the target for the confusion. With such a technique, it would be important that the fake profiles are representative of real users, in order to ensure that the confusion falls on other users rather than the correct one.

List of Figures

2.1	Inference of sensitive information and social relationships between users	19
2.2	Threat of mobility prediction Illustrated with a Mobility Markov Chain	20
2.3	Illustration of re-identification attacks	22
2.4	Different use case scenarios for the usage of LPPMs	27
2.5	Illustration of various types of mobility trace and record alteration . .	29
3.1	Re-identification attacks process from collecting phase to re-identification phase	40
3.2	From mobility trace to heat map	41
3.3	Performance of re-identification attacks on single output policy	51
3.4	Performance of LPPMs against re-identification attacks (single output policy)	53
3.5	Performance of LPPMs	54
3.6	Performance of the Threshold-Based Policy on Average Precision and Average False-Positive Rate Compared to the Theoretical Bound of the Top-k Policy	56

3.7	Effect of the number of users in the system to the user re-identification rate with single output policy (Cabspotting dataset)	57
3.8	Effect of the Proportion of the user mobility captured on the user re-identification rate with single output policy (Cabspotting datasets)	58
4.1	Overview of HMC	66
4.2	Heat Map alteration iterative process	66
4.3	Cell Number of Records Modification	68
4.4	Time Gaps constraining method	70
4.5	Comparison of HMC with competitors - Robustness against multiple attacks - CabSpotting & GeoLife datasets	79
4.6	Comparison of HMC with competitors - Robustness against multiple attacks - MDC & PrivaMov datasets	80
4.7	Detailed comparison of HMC with competitors - Anonymity set size against AP-Attack	81
4.8	Detailed comparison of HMC with competitors - Robustness against AP-Attack	82
4.9	Detailed comparison of HMC with competitors - Robustness against POI-Attack	83
4.10	Detailed comparison of HMC with competitors - Robustness against PIT-Attack	84
4.11	Detailed comparison of HMC with competitors - Multi-utility metrics evaluation	86
4.12	Detailed comparison of HMC with competitors - Area coverage utility metric	88

4.13 Detailed comparison of HMC with competitors - Spatial distortion utility metric	89
4.14 Detailed comparison of HMC with competitors - Spatial-temporal distortion utility metric	90
4.15 Detailed comparison of HMC with competitors - F-scores of surrounding POIs query utility metric	92
4.16 Comparison of HMC with competitors - Utility metric in terms of users' number of visits distortion – Cabspotting dataset	93
5.1 Re-identification rates of attacks in different scenarios with different datasets (80% training and 20% testing)	98
5.2 Re-identification of user by learning on short mobility traces	100
5.3 ILL-Attack compared to two state-of-the-art attacks in the session-based service scenario (30mn sessions – 80% training and 20% testing)	107
5.4 Resilience of ILL-Attack against LPPMs compared to two state-of-the-art attacks in the session-based service scenario (30mn sessions – 80% training and 20% testing)	108
5.5 ILL-Attack compared to two state-of-the-art attacks in the crowd-sensing scenario (24h data – 80% training and 20% testing)	110
5.6 Resilience of ILL-Attack against LPPMs compared to two state-of-the-art attacks in the crowd-sensing scenario (24h data – 80% training and 20% testing)	111
6.1 Illustration of how sub-traces have different subsets of LPPMs that protects them	116
6.2 Fine-grained protection mechanism with the sub-trace centric approach of LPPMs' protection	117

6.3	Comparison of <i>Hybrid-LPPM</i> with single LPPM competitors - Area coverage utility metric - 80 % for training and 20 % for testing	120
-----	---	-----

List of Tables

1.1	List of communications during the thesis	14
2.1	A Comparative between different use case scenarios of LPPMs	28
3.1	Example of attack policies	42
3.2	Description of datasets	49
4.1	Utility Measure Levels Description	85
5.1	List of features used by ILL-Attack	103
5.2	Description of the datasets used for the experiments on ILL-Attack .	105
6.1	Illustration of how different sub-traces have different re-identification results using a sample of sub-traces of the PrivaMov Dataset	117

Bibliography

- [1] Osman Abul, Francesco Bonchi, and Mirco Nanni. Anonymization of moving objects databases by clustering and perturbation. *Information Systems*, 35(8):884–910, 2010. ISSN 03064379. doi: 10.1016/j.is.2010.05.003. URL <http://dx.doi.org/10.1016/j.is.2010.05.003>. 8, 29, 38
- [2] Nadav Aharony, Wei Pan, Cory Ip, Inas Khayal, and Alex Pentland. Social fmri: Investigating and shaping social mechanisms in the real world. *Pervasive Mobile Computing*, 7(6):643–659, December 2011. 5
- [3] 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, WiSec 2018, Stockholm, Sweden, June 18-20, 2018*, pages 40–50, 2018. doi: 10.1145/3212480.3212483. URL <https://doi.org/10.1145/3212480.3212483>. 26, 34
- [4] 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>. 8, 25, 31, 38
- [5] Alastair R. Beresford and Frank Stajano. Location privacy in pervasive computing. *IEEE Pervasive Computing*, 2(1):46–55, 2003. doi: 10.1109/MPRV.2003.1186725. URL <https://doi.org/10.1109/MPRV.2003.1186725>. 34

- [6] Claudio Bettini, X Sean Wang, and Sushil Jajodia. Protecting Privacy Against Location-based Personal Identification. In *Proceedings of the Second VDLB International Conference on Secure Data Management, SDM'05*, pages 185–199, Berlin, Heidelberg, 2005. Springer-Verlag. ISBN 3-540-28798-1, 978-3-540-28798-8. doi: 10.1007/11552338_13. URL http://dx.doi.org/10.1007/11552338_13. 25
- [7] 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. *Wpes*, pages 179–188, 2013. ISSN 15437221. doi: 10.1145/2517840.2517842. 20
- [8] Vincent Bindschaedler and Reza Shokri. Synthesizing plausible privacy-preserving location traces. In *IEEE Symposium on Security and Privacy, SP 2016, San Jose, CA, USA, May 22-26, 2016*, pages 546–563, 2016. doi: 10.1109/SP.2016.39. URL <https://doi.org/10.1109/SP.2016.39>. 33, 35
- [9] Antoine Boutet, Sonia Ben Mokhtar, and Vincent Primault. Uniqueness assessment of human mobility on multi-sensor datasets. 2016. 49, 104
- [10] Leo Breiman. Technical note: Some properties of splitting criteria. *Machine Learning*, 24(1):41–47, 1996. doi: 10.1007/BF00117831. URL <https://doi.org/10.1007/BF00117831>. 105
- [11] Erik Buchmann, Klemens Böhm, Thorben Burghardt, and Stephan Kessler. Re-identification of smart meter data. *Personal and Ubiquitous Computing*, 17(4):653–662, 2013. doi: 10.1007/s00779-012-0513-6. URL <https://doi.org/10.1007/s00779-012-0513-6>. 21
- [12] 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*, 2018. URL http://wp.internetsociety.org/ndss/wp-content/uploads/sites/25/2018/02/ndss2018_06B-2_Caliskan_paper.pdf. 21
- [13] 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, 2017. doi: 10.1109/SRDS.2017.25. URL <https://doi.org/10.1109/SRDS.2017.25>. 8, 36, 74
- [14] Sung-hyuk Cha. Comprehensive Survey on Distance / Similarity Measures between Probability Density Functions. *International Journal of Mathematical Models and Methods in Applied Sciences*, 1(4):300–307, 2007. ISSN 14337347. doi: 10.1007/s00167-009-0884-z. URL <http://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.154.8446{%&}rep=rep1{%&}type=pdf>. 41, 42
- [15] Konstantinos Chatzikokolakis, Catuscia Palamidessi, and Marco Stronati. A predictive differentially-private mechanism for mobility traces. In *Privacy Enhancing Technologies - 14th International Symposium, PETS 2014, Amsterdam, The Netherlands, July 16-18, 2014. Proceedings*, pages 21–41, 2014. doi: 10.1007/978-3-319-08506-7_2. URL https://doi.org/10.1007/978-3-319-08506-7_2. 32
- [16] Konstantinos Chatzikokolakis, Catuscia Palamidessi, and Marco Stronati. Constructing elastic distinguishability metrics for location privacy. *PoPETs*, 2015 (2):156–170, 2015. doi: 10.1515/popets-2015-0023. URL <https://doi.org/10.1515/popets-2015-0023>. 31, 35
- [17] Konstantinos Chatzikokolakis, Ehab ElSalamouny, Catuscia Palamidessi, and Anna Pazii. Methods for location privacy: A comparative overview. *Foundations and Trends in Privacy and Security*, 1(4):199–257, 2017. doi: 10.1561/33000000017. URL <https://doi.org/10.1561/33000000017>. 31
- [18] Chi-Yin Chow, Mohamed F. Mokbel, and Xuan Liu. A peer-to-peer spatial cloaking algorithm for anonymous location-based service. In *14th ACM International Symposium on Geographic Information Systems, ACM-GIS 2006, November 10-11, 2006, Arlington, Virginia, USA, Proceedings*, pages 171–178, 2006. doi: 10.1145/1183471.1183500. URL <https://doi.org/10.1145/1183471.1183500>. 30
- [19] Daniel Cvrcek, Marek Kumpost, Vashek Matyas, and George Danezis. A study on the value of location privacy. In *Proceedings of the 2006 ACM Workshop on Privacy in the Electronic Society, WPES 2006, Alexandria, VA, USA*,

- October 30, 2006*, pages 109–118, 2006. doi: 10.1145/1179601.1179621. URL <https://doi.org/10.1145/1179601.1179621>. 6
- [20] 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, 2013. ISSN 2045-2322. 6, 22, 23
- [21] R Dingledine, N Mathewson, and P Syverson. Tor: The second-generation onion router, 2004. 50
- [22] Cynthia Dwork. Differential privacy: A survey of results. In *International Conference on Theory and Applications of Models of Computation*, pages 1–19. Springer, 2008. 8, 25, 31
- [23] Facebook. Facebook Places. URL <https://developers.facebook.com/docs/places/>. 19
- [24] Mustafa Amir Faisal, Alvaro A. Cárdenas, and Daisuke Mashima. How the quantity and quality of training data impacts re-identification of smart meter users? In *2015 IEEE International Conference on Smart Grid Communications, SmartGridComm 2015, Miami, FL, USA, November 2-5, 2015*, pages 31–36, 2015. doi: 10.1109/SmartGridComm.2015.7436272. URL <https://doi.org/10.1109/SmartGridComm.2015.7436272>. 21
- [25] Foursquare-Labs. Swarm, 2017. URL <https://www.swarmapp.com>. 4
- [26] Lorenzo Franceschi-Bicchierai. Redditor cracks anonymous data trove to pinpoint Muslim cab drivers. URL <https://mashable.com/2015/01/28/redditor-muslim-cab-drivers/>. 20
- [27] Julien Freudiger, Reza Shokri, and Jean-Pierre Hubaux. On the optimal placement of mix zones. In *Privacy Enhancing Technologies, 9th International Symposium, PETS 2009, Seattle, WA, USA, August 5-7, 2009. Proceedings*, pages 216–234, 2009. doi: 10.1007/978-3-642-03168-7_13. URL https://doi.org/10.1007/978-3-642-03168-7_13. 34
- [28] Sébastien Gambs, Marc-Olivier Killijian, and Miguel Nez Del Prado Cortez. Show Me How You Move and I Will Tell You Who You Are. *Transactions on Data Privacy*, 4:103–126, 2011. doi: 10.1145/1868470.1868479. 47, 63

- [29] 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. 5, 21
- [30] Sebastien Gambs, Marc-Olivier Killijian, and Miguel Nunez del Prado Cortez. De-anonymization Attack on Geolocated Data. *2013 12th IEEE International Conference on Trust, Security and Privacy in Computing and Communications*, pages 789–797, 2013. doi: 10.1109/TrustCom.2013.96. URL <http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=6680916>. 7, 23, 46, 47, 96
- [31] Bugra Gedik and Ling Liu. Location privacy in mobile systems: A personalized anonymization model. In *25th International Conference on Distributed Computing Systems (ICDCS 2005), 6-10 June 2005, Columbus, OH, USA*, pages 620–629, 2005. doi: 10.1109/ICDCS.2005.48. URL <https://doi.org/10.1109/ICDCS.2005.48>. 30
- [32] Pierre Geurts, Damien Ernst, and Louis Wehenkel. Extremely randomized trees. *Machine Learning*, 63(1):3–42, 2006. doi: 10.1007/s10994-006-6226-1. URL <https://doi.org/10.1007/s10994-006-6226-1>. 101, 103
- [33] Philippe Golle and Kurt Partridge. On the anonymity of home/work location pairs. In *Pervasive Computing, 7th International Conference, Pervasive 2009, Nara, Japan, May 11-14, 2009. Proceedings*, pages 390–397, 2009. doi: 10.1007/978-3-642-01516-8_26. URL https://doi.org/10.1007/978-3-642-01516-8_26. 7
- [34] Xiaowen Gong, Xu Chen, Kai Xing, Dong-Hoon Shin, Mengyuan Zhang, and Junshan Zhang. From social group utility maximization to personalized location privacy in mobile networks. *IEEE/ACM Trans. Netw.*, 25(3):1703–1716, 2017. doi: 10.1109/TNET.2017.2653102. URL <https://doi.org/10.1109/TNET.2017.2653102>. 34
- [35] Ian J. Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron C. Courville, and Yoshua Bengio. Generative adversarial nets. In *Advances in Neural Information Processing Systems 27: Annual Conference on Neural Information Processing Systems 2014, December 8-13 2014, Montreal, Quebec, Canada*, pages 2672–2680, 2014. URL <http://papers.nips.cc/paper/5423-generative-adversarial-nets>. 126

- [36] Google. Google Places. URL <https://developers.google.com/places/web-service/intro>. 19
- [37] Google. Google maps, 2017. URL <https://maps.google.com>. 4
- [38] Marco Gramaglia and Marco Fiore. Hiding mobile traffic fingerprints with GLOVE. In *Proceedings of the 11th ACM Conference on Emerging Networking Experiments and Technologies, CoNEXT 2015, Heidelberg, Germany, December 1-4, 2015*, pages 26:1–26:13, 2015. doi: 10.1145/2716281.2836111. URL <https://doi.org/10.1145/2716281.2836111>. 30
- [39] Marco Gramaglia, Marco Fiore, Alberto Tarable, and Albert Banchs. Preserving mobile subscriber privacy in open datasets of spatiotemporal trajectories. In *2017 IEEE Conference on Computer Communications, INFOCOM 2017, Atlanta, GA, USA, May 1-4, 2017*, pages 1–9. IEEE, 2017. ISBN 978-1-5090-5336-0. doi: 10.1109/INFOCOM.2017.8056979. URL <https://doi.org/10.1109/INFOCOM.2017.8056979>. 30, 35
- [40] Marco Gruteser and Dirk Grunwald. Anonymous usage of location-based services through spatial and temporal cloaking. In *Proceedings of the First International Conference on Mobile Systems, Applications, and Services, MobiSys 2003, San Francisco, CA, USA, May 5-8, 2003*, 2003. URL <http://www.usenix.org/events/mobisys03/tech/gruteser.html>. 30
- [41] Saikat Guha, Mudit Jain, and Venkata N. Padmanabhan. Koi: A location-privacy platform for smartphone apps. In *Proceedings of the 9th USENIX Symposium on Networked Systems Design and Implementation, NSDI 2012, San Jose, CA, USA, April 25-27, 2012*, pages 183–196, 2012. URL <https://www.usenix.org/conference/nsdi12/technical-sessions/presentation/guha>. 34
- [42] Gábor György Gulyás and Sándor Imre. Using identity separation against de-anonymization of social networks. *Trans. Data Privacy*, 8(2):113–140, 2015. URL <http://www.tdp.cat/issues11/tdp.a180a14.pdf>. 21
- [43] Nicolas Haderer, Romain Rouvoy, Christophe Ribeiro, and Lionel Seinturier. Apisense: Crowd-sensing made easy. *ERCIM News*, 93:28–29, 2013. 5

- [44] Ramaswamy Hariharan and Kentaro Toyama. Project Lachesis: Parsing and Modeling Location Histories. In Max J Egenhofer, Christian Freksa, and Harvey J Miller, editors, *Geographic Information Science: Third International Conference, GIScience 2004, Adelphi, MD, USA, October 20-23, 2004. Proceedings*, pages 106–124. Springer Berlin Heidelberg, Berlin, Heidelberg, 2004. ISBN 978-3-540-30231-5. doi: 10.1007/978-3-540-30231-5_8. URL http://dx.doi.org/10.1007/978-3-540-30231-5_8. 19, 46
- [45] 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 *Complex Networks and Their Applications VII - Volume 2 Proceedings The 7th International Conference on Complex Networks and Their Applications COMPLEX NETWORKS 2018*, pages 400–411, 2018. doi: 10.1007/978-3-030-05414-4_32. URL https://doi.org/10.1007/978-3-030-05414-4_32. 21
- [46] 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>. 29, 33
- [47] Kévin Huguenin, Igor Bilogrevic, Joana Soares Machado, Stefan Mihaila, Reza Shokri, Italo Dacosta, and Jean-Pierre Hubaux. A predictive model for user motivation and utility implications of privacy-protection mechanisms in location check-ins. *IEEE Trans. Mob. Comput.*, 17(4):760–774, 2018. doi: 10.1109/TMC.2017.2741958. URL <https://doi.org/10.1109/TMC.2017.2741958>. 5, 19
- [48] Sibren Isaacman, Richard A. Becker, Ramón Cáceres, Margaret Martonosi, James Rowland, Alexander Varshavsky, and Walter Willinger. Human mobility modeling at metropolitan scales. In *The 10th International Conference on Mobile Systems, Applications, and Services, MobiSys’12, Ambleside, United Kingdom - June 25 - 29, 2012*, pages 239–252, 2012. doi: 10.1145/2307636.2307659. URL <https://doi.org/10.1145/2307636.2307659>. 32
- [49] Zhipeng Jiang, Chao Zhao, Bin He, Yi Guan, and Jingchi Jiang. De-identification of medical records using conditional random fields and long

- short-term memory networks. *Journal of biomedical informatics*, 75:S43–S53, 2017. 21
- [50] Ryo Kato, Mayu Iwata, Takahiro Hara, Akiyoshi Suzuki, Xing Xie, Yuki Arase, and Shojiro Nishio. A dummy-based anonymization method based on user trajectory with pauses. In *SIGSPATIAL 2012 International Conference on Advances in Geographic Information Systems (formerly known as GIS), SIGSPATIAL'12, Redondo Beach, CA, USA, November 7-9, 2012*, pages 249–258, 2012. doi: 10.1145/2424321.2424354. URL <https://doi.org/10.1145/2424321.2424354>. 32
- [51] Carsten Kießler and Grant McKenzie. A geoprivacy manifesto. *Trans. GIS*, 22(1):3–19, 2018. doi: 10.1111/tgis.12305. URL <https://doi.org/10.1111/tgis.12305>. 4
- [52] John Krumm. Inference Attacks on Location Tracks. *Pervasive Computing*, 10 (Pervasive):127–143, 2007. ISSN 03029743. doi: 10.1007/978-3-540-72037-9_8. URL <http://www.springerlink.com/index/TG64551RW2716103.pdf> <http://research.microsoft.com/en-us/um/people/jckrumm/publications2007/inferenceattackrefined02distribute.pdf>. 7, 23
- [53] John Krumm and Dany Rouhana. Placer: semantic place labels from diary data. In *The 2013 ACM International Joint Conference on Pervasive and Ubiquitous Computing, UbiComp '13, Zurich, Switzerland, September 8-12, 2013*, pages 163–172, 2013. doi: 10.1145/2493432.2493504. URL <https://doi.org/10.1145/2493432.2493504>. 19
- [54] J K Laurila, Daniel Gatica-Perez, I Aad, Blom J., Olivier Bornet, Trinh-Minh-Tri Do, O Dousse, J Eberle, and M Miettinen. The Mobile Data Challenge: Big Data for Mobile Computing Research. In *Pervasive Computing*, 2012. 49, 104
- [55] Xinxin Liu, Han Zhao, Miao Pan, Hao Yue, Xiaolin Li, and Yuguang Fang. Traffic-aware multiple mix zone placement for protecting location privacy. In *Proceedings of the IEEE INFOCOM 2012, Orlando, FL, USA, March 25-30, 2012*, pages 972–980, 2012. doi: 10.1109/INFCOM.2012.6195848. URL <https://doi.org/10.1109/INFCOM.2012.6195848>. 34

- [56] Chris Y T Ma, David K Y Yau, Nung Kwan Yip, and Nageswara S V Rao. Privacy vulnerability of published anonymous mobility traces. *IEEE/ACM Transactions on Networking*, 21(3):720–733, 2013. ISSN 10636692. doi: 10.1109/TNET.2012.2208983. 23
- [57] Ashwin Machanavajjhala, Daniel Kifer, Johannes Gehrke, and Muthuramakrishnan Venkitasubramaniam. *L*-diversity: Privacy beyond *k*-anonymity. *TKDD*, 1(1):3, 2007. doi: 10.1145/1217299.1217302. URL <https://doi.org/10.1145/1217299.1217302>. 25
- [58] Mohamed Maouche. SFERA, 2017. URL <https://github.com/mmaouche-insa/SFERA/>. 48, 50
- [59] Mohamed Maouche, Sonia Ben Mokhtar, and Sara Bouchenak. Ap-attack: A novel user re-identification attack on mobility datasets. In *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, 2017. doi: 10.1145/3144457.3144494. URL <https://hal.archives-ouvertes.fr/hal-01785155/document>. 39, 96, 97, 105
- [60] Mohamed Maouche, Sonia Ben Mokhtar, and Sara Bouchenak. HMC: robust privacy protection of mobility data against multiple re-identification attacks. *IMWUT*, 2(3):124:1–124:25, 2018. doi: 10.1145/3264934. URL <https://hal.archives-ouvertes.fr/hal-01954041/document>.
- [61] Open Street Map. Amenity information description: <https://wiki.openstreetmap.org/wiki/key:amenity>, 2018. 77
- [62] Open Street Map. Download open street map dataset: https://wiki.openstreetmap.org/wiki/downloading_data, 2018. 31, 77
- [63] Sergio Mascetti, Dario Freni, Claudio Bettini, Xiaoyang Sean Wang, and Sushil Jajodia. Privacy in geo-social networks: proximity notification with untrusted service providers and curious buddies. *VLDB J.*, 20(4):541–566, 2011. doi: 10.1007/s00778-010-0213-7. URL <https://doi.org/10.1007/s00778-010-0213-7>. 8, 26, 34
- [64] Microsoft. Bing maps, 2017. URL <https://www.bing.com/maps>. 4

- [65] Darakhshan J. Mir, Sibren Isaacman, Ramón Cáceres, Margaret Martonosi, and Rebecca N. Wright. DP-WHERE: differentially private modeling of human mobility. In *Proceedings of the 2013 IEEE International Conference on Big Data, 6-9 October 2013, Santa Clara, CA, USA*, pages 580–588, 2013. doi: 10.1109/BigData.2013.6691626. URL <https://doi.org/10.1109/BigData.2013.6691626>. 32
- [66] Prashanth Mohan, Venkata N. Padmanabhan, and Ramachandran Ramjee. Ner-cell: Rich monitoring of road and traffic conditions using mobile smartphones. In *SenSys*, pages 323–336, 2008. 5
- [67] Mohamed F. Mokbel, Chi-Yin Chow, and Walid G. Aref. The new casper: Query processing for location services without compromising privacy. In *Proceedings of the 32nd International Conference on Very Large Data Bases, Seoul, Korea, September 12-15, 2006*, pages 763–774, 2006. URL <http://dl.acm.org/citation.cfm?id=1164193>. 30
- [68] Yoni De Mulder, George Danezis, Lejla Batina, and Bart Preneel. Identification via location-profiling in GSM networks. In *Proceedings of the 2008 ACM Workshop on Privacy in the Electronic Society, WPES 2008, Alexandria, VA, USA, October 27, 2008*, pages 23–32, 2008. doi: 10.1145/1456403.1456409. URL <https://doi.org/10.1145/1456403.1456409>. 23
- [69] Min Mun, Sasank Reddy, Katie Shilton, Nathan Yau, Jeff Burke, Deborah Estrin, Mark Hansen, Eric Howard, Ruth West, and Péter Boda. Peir, the personal environmental impact report, as a platform for participatory sensing systems research. In *MobiSys*, pages 55–68, 2009. 5
- [70] F.M. Naini, J. Unnikrishnan, P. Thiran, and M. Vetterli. Where You Are Is Who You Are: User Identification by Matching Statistics. *IEEE Transactions on Information Forensics and Security*, 11(2):358–372, 2016. ISSN 1556-6013. doi: 10.1109/TIFS.2015.2498131. URL [http://ieeexplore.ieee.org/xpl/articleDetails.jsp?arnumber=7321027&isnumber=7349121&tp={&}arnumber=7321027{&}isnumber=7349121{&}5Cnhttp://ieeexplore.ieee.org/xpl/articleDetails.jsp?arnumber=7321027{&}filter=AND\(p{&}IS{&}Number:7349121\)](http://ieeexplore.ieee.org/xpl/articleDetails.jsp?arnumber=7321027&isnumber=7349121&tp={&}arnumber=7321027{&}isnumber=7349121{&}5Cnhttp://ieeexplore.ieee.org/xpl/articleDetails.jsp?arnumber=7321027{&}filter=AND(p{&}IS{&}Number:7349121)). 23
- [71] 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, 2008. doi: 10.1109/SP.2008.33. URL <https://doi.org/10.1109/SP.2008.33>. 21
- [72] Niantic. Pokemon Go, 2017. URL <http://www.pokemongo.com>. 4
- [73] Niantic. Harry Potter Wizards Unite, 2019. URL <https://www.harrypotterwizardsunite.com/>. 4
- [74] Anastasios Noulas, Salvatore Scellato, Neal Lathia, and Cecilia Mascolo. Mining user mobility features for next place prediction in location-based services. In *12th IEEE International Conference on Data Mining, ICDM 2012, Brussels, Belgium, December 10-13, 2012*, pages 1038–1043, 2012. doi: 10.1109/ICDM.2012.113. URL <https://doi.org/10.1109/ICDM.2012.113>. 20
- [75] Randal S. Olson, Nathan Bartley, Ryan J. Urbanowicz, and Jason H. Moore. Evaluation of a tree-based pipeline optimization tool for automating data science. In *Proceedings of the Genetic and Evolutionary Computation Conference 2016, GECCO '16*, pages 485–492, New York, NY, USA, 2016. ACM. ISBN 978-1-4503-4206-3. doi: 10.1145/2908812.2908918. URL <http://doi.acm.org/10.1145/2908812.2908918>. 104, 105
- [76] Balaji Palanisamy and Ling Liu. Mobimix: Protecting location privacy with mix-zones over road networks. In *Proceedings of the 27th International Conference on Data Engineering, ICDE 2011, April 11-16, 2011, Hannover, Germany*, pages 494–505, 2011. doi: 10.1109/ICDE.2011.5767898. URL <https://doi.org/10.1109/ICDE.2011.5767898>. 34
- [77] Sai Teja Peddinti and Nitesh Saxena. On the limitations of query obfuscation techniques for location privacy. In *UbiComp 2011: Ubiquitous Computing, 13th International Conference, UbiComp 2011, Beijing, China, September 17-21, 2011, Proceedings*, pages 187–196, 2011. doi: 10.1145/2030112.2030139. URL <https://doi.org/10.1145/2030112.2030139>. 33
- [78] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay. Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, 12: 2825–2830, 2011. 105

- [79] Albin Petit, Thomas Cerqueus, Antoine Boutet, Sonia Ben Mokhtar, David Coquil, Lionel Brunie, and Harald Kosch. Simattack: private web search under fire. *J. Internet Services and Applications*, 7(1):2:1–2:17, 2016. doi: 10.1186/s13174-016-0044-x. URL <https://doi.org/10.1186/s13174-016-0044-x>. 21
- [80] Michal Piorkowski, Natasa Sarafijanovic-djukic, and Matthias Grossglauser. CRAW- DAD data set epfl/mobility (v. 2009-02-24), 2009. URL <http://crawdad.cs.dartmouth.edu/epfl/mobility>. 49, 104
- [81] V. Primault, A. Boutet, S. B. Mokhtar, and L. Brunie. The long road to computational location privacy: A survey. *IEEE Communications Surveys Tutorials*, pages 1–1, 2018. ISSN 1553-877X. doi: 10.1109/COMST.2018.2873950. 19
- [82] Vincent Primault, Sonia Ben Mokhtar, Cédric Lauradoux, and Lionel Brunie. Differentially Private Location Privacy in Practice. *Most'14*, (October), 2014. 23, 46, 48, 63, 96, 97
- [83] 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, 2015. doi: 10.1109/Trustcom.2015.417. URL <https://doi.org/10.1109/Trustcom.2015.417>. 7, 8, 26, 29, 35, 36, 38, 68, 72, 75
- [84] 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, 2018. doi: 10.1109/ICDCS.2018.00091. URL <https://hal.archives-ouvertes.fr/hal-01784557v2/document>.
- [85] Jean Louis Raisaro, Florian Tramèr, Zhanglong Ji, Diyue Bu, Yongan Zhao, Knox Carey, David Lloyd, Heidi Sofia, Dixie Baker, Paul Flicek, Suyash Shringarpure, Carlos 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. *Journal of the American Medical Informatics Association*, 24(4):799–805, 02 2017. ISSN 1527-974X. doi: 10.1093/jamia/ocw167. URL <https://doi.org/10.1093/jamia/ocw167>. 21
- [86] 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, 2014. doi: 10.1109/PerCom.2014.6813960. URL <https://doi.org/10.1109/PerCom.2014.6813960>. 26, 35
- [87] Marco Romanelli, Catuscia Palamidessi, and Konstantinos Chatzikokolakis. Generating optimal privacy-protection mechanisms via machine learning. *CoRR*, abs/1904.01059, 2019. URL <http://arxiv.org/abs/1904.01059>. 126
- [88] S2Geometry. Computational geometry and spatial indexing on the sphere <http://s2geometry.io/>. 105
- [89] Adam Sadilek and John Krumm. Far out: Predicting long-term human mobility. In *Proceedings of the Twenty-Sixth AAAI Conference on Artificial Intelligence, July 22-26, 2012, Toronto, Ontario, Canada.*, 2012. URL <http://www.aaai.org/ocs/index.php/AAAI/AAAI12/paper/view/4845>. 21
- [90] Pierangela Samarati and Latanya Sweeney. Generalizing data to provide anonymity when disclosing information (abstract). In *Proceedings of the Seventeenth ACM SIGACT-SIGMOD-SIGART Symposium on Principles of Database Systems, June 1-3, 1998, Seattle, Washington, USA*, page 188, 1998. doi: 10.1145/275487.275508. URL <https://doi.org/10.1145/275487.275508>. 25
- [91] Pierangela Samarati and Latanya Sweeney. Generalizing Data to Provide Anonymity when Disclosing Information. In *Proceedings of the Seventeenth ACM SIGACT-SIGMOD-SIGART Symposium on Principles of Database Systems, PODS '98*, pages 188—, New York, NY, USA, 1998. ACM. ISBN 0-89791-996-3. doi: 10.1145/275487.275508. URL <http://doi.acm.org/10.1145/275487.275508>. 8
- [92] Felix Schuster, Manuel Costa, Cédric Fournet, Christos Gkantsidis, Marcus Peinado, Gloria Mainar-Ruiz, and Mark Russinovich. VC3: trustworthy data analytics in the cloud using SGX. In *2015 IEEE Symposium on Security and*

- Privacy, SP 2015, San Jose, CA, USA, May 17-21, 2015*, pages 38–54, 2015. doi: 10.1109/SP.2015.10. URL <https://doi.org/10.1109/SP.2015.10>. 114
- [93] Pravin Shankar, Vinod Ganapathy, and Liviu Iftode. Privately querying location-based services with sybilquery. In *UbiComp 2009: Ubiquitous Computing, 11th International Conference, UbiComp 2009, Orlando, Florida, USA, September 30 - October 3, 2009, Proceedings*, pages 31–40, 2009. doi: 10.1145/1620545.1620550. URL <https://doi.org/10.1145/1620545.1620550>. 33
- [94] Mudhakar Srivatsa and Mike Hicks. Deanonymizing Mobility Traces : Using Social Networks as a Side-Channel. *Proceedings of the 2012 ACM conference on Computer and communications security (CCS)*, pages 628–637, 2012. ISSN 15437221. doi: 10.1145/2382196.2382262. 23
- [95] Jacopo Staiano, Nuria Oliver, Bruno Lepri, Rodrigo de Oliveira, Michele Caraviello, and Nicu Sebe. Money walks: a human-centric study on the economics of personal mobile data. In *The 2014 ACM Conference on Ubiquitous Computing, UbiComp '14, Seattle, WA, USA, September 13-17, 2014*, pages 583–594, 2014. doi: 10.1145/2632048.2632074. URL <https://doi.org/10.1145/2632048.2632074>. 5
- [96] Latanya Sweeney. k-anonymity: A model for protecting privacy. *International Journal of Uncertainty, Fuzziness and Knowledge-Based Systems*, 10(5):557–570, 2002. doi: 10.1142/S0218488502001648. URL <https://doi.org/10.1142/S0218488502001648>. 25
- [97] 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, 2018. doi: 10.1109/SP.2018.00008. URL <https://doi.org/10.1109/SP.2018.00008>. 21
- [98] Chen Wang, Chuyu Wang, Yingying Chen, Lei Xie, and Sanglu Lu. Smartphone privacy leakage of social relationships and demographics from surrounding access points. In *37th IEEE International Conference on Distributed Computing Systems, ICDCS 2017, Atlanta, GA, USA, June 5-8, 2017*, pages 678–688, 2017. doi: 10.1109/ICDCS.2017.139. URL <https://doi.org/10.1109/ICDCS.2017.139>. 5, 20

- [99] Ningfei Wang, Shouling Ji, and Ting Wang. Integration of static and dynamic code stylometry analysis for programmer de-anonymization. In *Proceedings of the 11th ACM Workshop on Artificial Intelligence and Security, CCS 2018, Toronto, ON, Canada, October 19, 2018*, pages 74–84, 2018. doi: 10.1145/3270101.3270110. URL <https://doi.org/10.1145/3270101.3270110>. 21
- [100] Marius Wernke, Pavel Skvortsov, Frank Dürri, and Kurt Rothermel. A classification of location privacy attacks and approaches. *Personal and Ubiquitous Computing*, 18(1):163–175, 2014. doi: 10.1007/s00779-012-0633-z. URL <https://doi.org/10.1007/s00779-012-0633-z>. 18
- [101] Yonghui Xiao and Li Xiong. Protecting locations with differential privacy under temporal correlations. In *Proceedings of the 22nd ACM SIGSAC Conference on Computer and Communications Security, Denver, CO, USA, October 12-16, 2015*, pages 1298–1309, 2015. doi: 10.1145/2810103.2813640. URL <https://doi.org/10.1145/2810103.2813640>. 32
- [102] Yonghui Xiao, Li Xiong, Si Zhang, and Yang Cao. Loclok: Location cloaking with differential privacy via hidden markov model. *PVLDB*, 10(12):1901–1904, 2017. doi: 10.14778/3137765.3137804. URL <http://www.vldb.org/pvldb/vol10/p1901-xiao.pdf>. 32
- [103] Toby Xu and Ying Cai. Feeling-based location privacy protection for location-based services. In *Proceedings of the 2009 ACM Conference on Computer and Communications Security, CCS 2009, Chicago, Illinois, USA, November 9-13, 2009*, pages 348–357, 2009. doi: 10.1145/1653662.1653704. URL <https://doi.org/10.1145/1653662.1653704>. 26, 30
- [104] Mingqiang Xue, Panos Kalnis, and Hung Keng Pung. Location diversity: Enhanced privacy protection in location based services. In *Location and Context Awareness, 4th International Symposium, LoCA 2009, Tokyo, Japan, May 7-8, 2009, Proceedings*, pages 70–87, 2009. doi: 10.1007/978-3-642-01721-6_5. URL https://doi.org/10.1007/978-3-642-01721-6_5. 25
- [105] Tun-Hao You, Wen-Chih Peng, and Wang-Chien Lee. Protecting moving trajectories with dummies. In *8th International Conference on Mobile Data Management (MDM 2007), Mannheim, Germany, May 7-11, 2007*, pages 278–282, 2007. doi: 10.1109/MDM.2007.58. URL <https://doi.org/10.1109/MDM.2007.58>. 32

-
- [106] Yu Zheng, Xing Xie, and Wei-Ying Ma. GeoLife: A Collaborative Social Networking Service among User, location and trajectory. *IEEE Data(base) Engineering Bulletin*, 2010. 49, 104
- [107] 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*, GIS '04, pages 266–273, New York, NY, USA, 2004. ACM. ISBN 1-58113-979-9. doi: 10.1145/1032222.1032261. URL <http://doi.acm.org/10.1145/1032222.1032261>. 19, 46



FOLIO ADMINISTRATIF

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

NOM : MAOUCHE

DATE de SOUTENANCE : 06/11/2019

Prénoms : Mohamed

TITRE : Protection against Re-identification Attacks in Location Privacy

NATURE : Doctorat

Numéro d'ordre : 2019LYSEI089

Ecole doctorale : InfoMaths (ED 5120)

Spécialité : Informatique

RESUME :

De nos jours, avec la large propagation de différents appareils mobiles, de nombreux capteurs accompagnent des utilisateurs. Ces capteurs peuvent servir à collecter des données de mobilité qui sont utiles pour de urbanistes ou des chercheurs. Cependant, l'exploitation de ces données soulèvent de nombreuses menaces quant à la préservation de la vie privée des utilisateurs. En effet, des informations sensibles tel que le lieu domicile, le lieu de travail ou même les croyances religieuses peuvent être inférées de ces données. Durant la dernière décennie, des mécanismes de protections appelées "Location Privacy Protection Mechanisms (LPPM)" ont été proposés. Ils imposent des garanties sur les données (e.g., k-anonymity ou differential privacy), obfusquent les informations sensibles (e.g., efface les points d'intérêt) ou sont une contremesure à des attaques particulières. Nous portons notre attention à la ré-identification qui est un risque précis lié à la préservation de la vie privée dans les données de mobilité. Il consiste en a un attaquant qui dès lors qu'il reçoit une trace de mobilité anonymisée, il cherche à retrouver l'identifiant de son propriétaire en la rattachant à un passif de traces non-anonymisées des utilisateurs du système. Dans ce cadre, nous proposons tout d'abords des attaques de ré-identification AP-Attack et ILL-Attack servant à mettre en exergue les vulnérabilités des mécanismes de protections de l'état de l'art et de quantifier leur efficacité. Nous proposons aussi un nouveau mécanisme de protection HMC qui utilise des heat maps afin de guider la transformation du comportement d'un individu pour qu'il ne ressemble plus au soi du passée mais à un autre utilisateur, le préservant ainsi de la ré-identification. Cette modification de la trace de mobilité est contrainte par des mesures d'utilité des données afin de minimiser la qualité de service ou les conclusions que l'on peut tirer à l'aide de ces données.

MOTS-CLÉS : Vie privée, données de mobilité, attaques de ré-identification, mécanismes de protections, utilité des données.

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

Directeur de thèse:

Bouchenak, Sara
Ben Mokhtar Sonia

Professeur des Universités, INSA-Lyon
Directrice de Recherche, CNRS

Directrice de thèse
Co-directrice de thèse

Président de jury :

Composition du jury :

Fiore, Marco
Palamidessi, Catuscia
Castelluccia, Claude
Petit, Jean-Marc

Habilité à Diriger des Recherches, CNR-IEIT
Directrice de Recherche, INRIA
Directeur de Recherche, INRIA
Professeur des Universités, INSA-Lyon

Rapporteur
Rapporteuse
Examineur
Examineur