



Labs for prototyping future mobility data sharing solutions in the cloud

D2.3 State of the art on Mobility and Transport data protection technologies

23/01/2023

Author(s): Alberto BLANCO-JUSTICIA (URV), Emre BAYAMLIOĞLU (KUL), Alik
BENMAYOR (KUL), Benet MANZANARES SALOR (URV)



MobiDataLab is funded by the EU
under the H2020 Research and
Innovation Programme (grant
agreement No 101006879).

Summary sheet

Deliverable Number	D2.3
Deliverable Name	State of the Art on Mobility and Transport data protection technologies
Full Project Title	MobiDataLab, Labs for prototyping future Mobility Data sharing cloud solutions
Responsible Author(s)	Alberto BLANCO-JUSTICIA (URV)
Contributing Partner(s)	HERE, HOVE, KUL
Peer Review	AKKA, CNR
Contractual Delivery Date	31-07-2021
Actual Delivery Date	29-07-2021
Status	Final
Dissemination level	Public
Version	V1.0
No. of Pages	75
WP/Task related to the deliverable	WP2/T2.2
WP/Task responsible	AKKA/URV
Document ID	MobiDataLab-D2.3-SoATransportDataProtTechno-v1.0
Abstract	This document describes the legal provisions and risks associated with the collection of personal mobility and transportation data, along with design strategies and technologies to reduce or eliminate such risks.

Legal Disclaimer

MOBIDATALAB (Grant Agreement No 101006879) is a Research and Innovation Actions project funded by the EU Framework Programme for Research and Innovation Horizon 2020. This document contains information on MOBIDATALAB core activities, findings, and outcomes. The content of this publication is the sole responsibility of the MOBIDATALAB consortium and cannot be considered to reflect the views of the European Commission.

Project partners

Organisation	Country	Abbreviation
AKKA I&S	France	AKKA
CONSORZIO INTERUNIVERSITARIO PER L'OTTIMIZZAZIONE E LA RICERCA OPERATIVA	Italy	ICOOR
AETHON SYMVOULI MICHANIKI MONOPROSOPI IKE	Greece	AETHON
CONSIGLIO NAZIONALE DELLE RICERCHE	Italy	CNR
HOVE	France	HOVE
HERE GLOBAL B.V.	Netherlands	HERE
KATHOLIEKE UNIVERSITEIT LEUVEN	Belgium	KUL
UNIVERSITAT ROVIRA I VIRGILI	Spain	URV
POLIS - PROMOTION OF OPERATIONAL LINKS WITH INTEGRATED SERVICES	Belgium	POLIS
F6S NETWORK IRELAND LIMITED	Ireland	F6S

Document history

Version	Date	Organisation	Main area of changes	Comments
0.1	01/03/2021	URV, KUL, HOVE, HERE	Table of contents	
0.5	04/06/2021	URV, KUL	Sections 2, 4, 5, 6, 7	
0.6	09/06/2021	URV	Changed to new document template	
0.7	13/06/2021	URV	Sections 1, 8	
0.8	21/06/2021	URV	All	Rework
1.0	29/07/2021	AKKA		Quality check and submission

Executive Summary

The main goal of this task and deliverable is to produce a theoretically and practically oriented overview of available privacy-by-design strategies, privacy risk assessment tools and privacy-preserving technologies that will allow us to develop responsible, ethically aligned and GDPR-compliant tools for the storage, sharing and analysis of mobility and transport data, subject to GDPR principles, such as purpose limitation. Different privacy models and anonymisation techniques and methods will be analysed, considering their applicability to the project's use cases with regards to i) their source of inputs (e.g., data from fixed sensors, from user devices, from publicly available datasets, etc.), and ii) the kind of scenario in which they can be applied (e.g., analysis of raw mobility data, data sharing or data release). Anonymisation methods with the potential to achieve the privacy goals of the project will be studied.

Methods that allow for the consolidation of anonymised data sets will be of particular interest, since the study of mobility and transport data might require the fusion and enrichment of data using information from several sources.

Table of contents

1. INTRODUCTION.....	8
1.1. PROJECT OVERVIEW.....	8
1.2. PURPOSE OF THIS DELIVERABLE.....	8
1.3. STRUCTURE OF THE DELIVERABLE.....	9
2. THE PRIVACY AND DATA PROTECTION REGULATORY FRAMEWORK	10
2.1. PRIVACY AND PERSONAL DATA PROTECTION UNDER THE ECHR AND CFR.....	11
2.2. DATA PROTECTION UNDER THE GDPR	12
3. SOURCES AND USES OF LOCATION DATA.....	17
4. GENERAL DATA PRIVACY MEASURES	20
4.1. DATA PROTECTION BY DESIGN	21
4.2. TOOLS AND TECHNIQUES.....	22
5. PRIVACY OF LOCATION-BASED SERVICES	38
5.1. UNIQUE CHARACTERISTICS.....	38
5.2. RISKS	39
5.3. PRIVACY METHODS	40
6. PRIVACY OF TRAJECTORY MICRODATA SHARING AND RELEASE.....	58
6.1. UNIQUE CHARACTERISTICS.....	58
6.2. RISKS	59
6.3. PRIVACY METHODS	61
7. PRIVACY OF AGGREGATED MOBILITY DATA	70
7.1. UNIQUE CHARACTERISTICS.....	70
7.2. RISKS	70
7.3. PRIVACY METHODS	72
8. CONCLUSIONS.....	73
9. ANNEXES	74

List of figures

Figure 1: Example of the Group-based approach with six users. The two groups are coloured green and blue respectively. Each user is represented by a circle with the pseudonym inside. Rectangles are used to depict the cloaking regions.	42
Figure 2: Example of the Distortion-based method with six users. The two groups are differentiated by colour. Each user is represented by a circle with a character as pseudonym and an arrow depicting their expected movement direction. Rectangular regions are used to illustrate the cloaking areas.....	43
Figure 3: Example of the Mix-zones method with three users. Each pseudonym's trajectory has a different colour, a start and end position defined by a circle with the pseudonym inside and an arrow defining the trajectory. The Mix-zone region is defined by a white rectangle. The user's pseudonyms have the following changes: A→F, B→D, C→E.	44
Figure 4: Example of SwapMob with three users. Each pseudonym's trajectory has a different colour, an initial and final position defined by a circle with the pseudonym inside and an arrow which defining the path.....	46

- Figure 5: Example of P2P spatial cloaking for the user A. On the left side, the P2P network is depicted, using edges to represent connections between nodes. Blue nodes are one hop away from A, and those two hops away are coloured purple. The right-hand side of the figure represents the resulting cloaking region of A for a k of four. 48
- Figure 6: Example of CliqueCloak with three users at time t1 and four users at time t2. Users are represented as coloured circles with the character of the user's pseudonym. The maximum spatial tolerance is defined by squares centred in the corresponding user. The current constraint graph is shown on the right. The user's desired k is displayed next to the respective node of the constraint graph. The resulting cloaking region is represented as a grey rectangle. 52
- Figure 8: GLOVE Spatiotemporal generalisation procedure. 63
- Figure 9: NWA uncertain trajectory and possible motion curve in red. 65

List of tables

- Table 1: Location data required in the use cases. Includes whether the data are personal information or not, and which Section(s) in this document deal(s) with the protection of such data. 18
- Table 2: Example microdata table: The Name attribute is an Identifier, Age and Zipcode are quasi-identifiers, while the Diagnostic attribute is a Confidential attribute. 32
- Table 3: Main characteristics of the privacy methods at a glance. 56
- Table 4: Overview of strengths and weaknesses of techniques assessed by Article 29 Working Party in Opinion 5/2014. 74

Abbreviations and acronyms

Abbreviation	Meaning
CFR	European Charter of Fundamental Rights
CJEU	Court of Justice of the EU
DHT	Distributed Hash Table
ECHR	European Convention of Human Rights
ETA	Estimated Time of Arrival
EU	European Union
GDPR	General Data Protection Regulation

GPS	Global Positioning System
IFS	Interpolation First and Sampling
INE	Instituto Nacional de Estadística (Spanish National Institute of Statistics)
IP	Internet Protocol
LBQID	Location-Based Quasi-Identifier
LBS	Location-Based Service(s)
MS	Member State(s)
NIST	National Institute of Standards and Technology
NWA	Never Walk Alone
P2P	Peer-to-Peer
PoI	Point of Interest
RFID	Radio-frequency Identification
SFI	Sampling First and Interpolation
TFEU	Treaty for the Functioning of the European Union
TTP	Trusted Third Party

1. Introduction

1.1. Project overview

There has been an explosion of mobility services and data sharing in recent years. Building on this, the EU-funded MobiDataLab project works to foster the sharing of data amongst transport authorities, operators, and other mobility stakeholders in Europe. MobiDataLab develops knowledge as well as a cloud solution aimed at easing the sharing of data. Specifically, the project is based on a continuous co-development of knowledge and technical solutions. It collects and analyses the advice and recommendations of experts and supporting cities, regions, clusters, and associations. These actions are assisted by the incremental construction of a cross-thematic knowledge base and a cloud-based service platform, which will improve access and usage of data sharing resources.

We thank Jordi Soria-Comas from the Catalan Data Protection Authority, Leo Frachet, and Tu-Tho Thai from MobilityData for their insightful comments during the preparation of this document.

1.2. Purpose of this deliverable

This deliverable offers an overview of the privacy risks related to mobility and transportation data and provides a description to several techniques in the literature aiming at limiting or eliminating such risks.

Mobility data in its simpler form are data about individuals that include their locations at specific times. Sources of real-time raw individual location data include, but are not limited to, cell towers, Wi-Fi access points, radio-frequency identification (RFID) tag readers, location-based services, or credit card payments. Historical location data, in form of data sets in which each of the records corresponds to an individual and includes her location data for some time periods are referred to as trajectory microdata sets. Such trajectory microdata sets are often of interest to transport authorities, operators, and other stakeholders to evaluate and improve their services, the state of the traffic, etc. and thus are often publicly released or shared. Sharing of mobility data is occasionally shared as aggregates (e.g., heat maps) instead of at an individual level.

Whichever the form of these mobility data, they all share some statistical characteristics that make their sharing a potential privacy risk. Mobility data are highly *unique* and *regular*. *Unicity* refers to the data of different individuals to be easily differentiable, particularly at some specific locations. The starting and ending locations of users' trajectories are often their home and work locations which, again are highly unique and can lead to reidentification. Studies show that user full trajectories can be uniquely recovered with the knowledge of only two locations. The *regularity* of trajectories means that for single individuals, their data follows periodic patterns. Namely, individuals tend to follow the same trajectories during workdays—home to work and back to home.

This deliverable introduces and describes the legal and regulatory implications of the collection and processing of mobility data in the context of the General Data Protection Regulation (GDPR), and then studies the risks of reidentification for location-based services, trajectory microdata sharing and

release, and aggregated data sharing and release, along with techniques in the literature that aim at reducing such risks.

The objective of this deliverable is to help in the requirement elicitation for the implementation of the project's use cases, along with the prioritisation of data protection techniques more suitable to deal with the particularities of the use cases.

1.3. Structure of the deliverable

This deliverable is organised as follows. Section 2 introduces the legal and regulatory framework for privacy and data protection in the context of mobility data. Section 3 briefly describes the kinds of mobility data and the processing operations covered by the MobiDataLab use cases. Section 4 introduces general data protection strategies and technologies, such as privacy-by-design strategies, encryption, and anonymisation techniques. Section 5 describes the characteristics of location-based services and several techniques to protect said location data. Section 6 does the same for trajectory microdata sharing and release, and Section 7 deals with aggregate data. We provide conclusions in Section 8.

2. The Privacy and Data Protection Regulatory Framework

MobiDataLab is expected to collect large amounts of mobility and transport data that will include personal data (e.g., passenger, vehicle, travel data). Any collection, processing, and sharing of personal data need to be performed following international and European legislation and any interference with the fundamental rights, such as the right to privacy and personal data protection, must be avoided or at least, be necessary and proportionate.

Both rights are protected under the following frameworks:

- The European Convention on Human Rights (ECHR)¹: this is an international agreement in the framework of the Council of Europe, an organisation that includes member states (MS) of the European Union (EU) and some non-EU MS. The ECHR was adopted in 1950 and entered into force in 1953. Signatory states must respect the rights stipulated in the ECHR when exercising any activity or power.
- The Treaty for the Functioning of the European Union (TFEU), including the European Charter of Fundamental Rights (CFR): the CFR includes all the personal, civic, political, economic, and social rights enjoyed by individuals within the EU. It became legally binding as EU primary law² in 2009. The provisions of the CFR are addressed to EU institutions and bodies, obliging them to respect the rights listed therein when fulfilling their duties. The CFR also binds EU MS when implementing EU law.

As far as EU law is concerned, the General Data Protection Regulation (GDPR)³ adopted in 2016 is currently the legislation that establishes a detailed and comprehensive data protection system in the EU.

¹ The European Court of Human Rights has competence to adjudicate upon complaints by individuals on alleged breaches of human rights by signatory states.

² EU law is composed of primary and secondary EU law. Primary EU law is the supreme source of law in the EU. It sets out the distribution of powers and responsibilities between the EU and EU countries and provides the legal context within which EU institutions formulate and implement policies. The treaties, namely the Treaty on European Union (TEU) and the TFEU form part of primary EU law. Secondary EU law is the body of law adopted by the EU institutions based on the principles and objectives set out in the Treaties. Examples include Regulations, Directives, and decisions of the EU.

³ Regulation (EU) 2016/679 of the European Parliament and the Council of 27 April 2016 on the protection of natural persons regarding the processing of personal data and on the free movement of such data, and repealing Directive 95/46/EC.

This section will, first, briefly analyse the scope of the rights under the ECHR and the CFR. The analysis touches upon both the CFR and the ECHR as the scope of the rights protected is similar⁴ and both the Court of Justice of the EU (CJEU) and the European Court of Human Rights (ECtHR) refer to the jurisprudence of each other. Then, we will focus mainly on the relevant provisions of the GDPR, as the text most relevant for MobiDataLab.

2.1. Privacy and personal data protection under the ECHR and CFR

Articles 8 ECHR and 7 CFR respectively, protect the right to private life, home, and communications.⁵ According to the case law of the ECHR, the concept of private life extends to aspects relating to personal identity, such as a person's name or a photo. It should be considered as a broad term, not susceptible to exhaustive definition. It covers the physical and psychological integrity of a person. It can embrace multiple aspects of a person's identity, such as gender identification or name, or elements relating to a person's right to their image.⁶ It covers personal information which individuals can legitimately expect should not be published without their consent.⁷

A person's "private life" may equally cover situations outside a person's home or private premises. However, that may not be the case when people *knowingly* or *intentionally* involve themselves in activities that are or may be recorded in a public manner (e.g., using public transport where surveillance cameras exist). But private life considerations may arise once any systematic or permanent record comes into existence of such material from the public domain.⁸

Article 8 is the main provision for the protection of personal data under the ECHR. The right is further protected in Article 16 of the TFEU and Article 8 of the CFR, which is based on Article 8 ECHR.⁹ The first paragraph of Article 8 CFR provides that "[e]veryone has the right to the protection of

⁴ According to Article 52(3) of the CFR states that the meaning and scope of a Charter right corresponding to a right in the ECHR shall be the same as those laid down by the ECHR.

⁵ Article 8 ECHR provides that "Everyone has the right to respect for his private and family life, his home and his correspondence". Article 7 CFR provides that "Everyone has the right to respect for his or her private and family life, home and communications". The explanatory note on Article 7 provides that the rights guaranteed in Article 7 CFR correspond to those guaranteed by Article 8 of the ECHR. To take account of developments in technology, the word "correspondence" has been replaced by "communications".

⁶ European Court of Human Rights, Guide on Article 8 of the European Convention on Human Rights, updated on 31 December 2020: https://www.echr.coe.int/Documents/Guide_Art_8_ENG.pdf.

⁷ Case *Axel Springer AG v. Germany* [GC], Application no. 39954/08, para. 83.

⁸ Case *P.G. and J.H. v. the United Kingdom*, Application no. 44787/98, para. 57.

⁹ See explanations relating to the Charter of Fundamental Rights, OJ C303/20.

*personal data concerning him or her*¹⁰. The roots of the right to personal data protection lie in the right to privacy. Although the two rights are closely linked, they should not be seen as the same right. A separate right to data protection was included in the CFR to provide a conclusive system of checks and balances which ensures lawful processing of personal data.¹¹ Indeed, the second paragraph of Article 8 sets out that “[s]uch data must be processed fairly for specified purposes and on the basis of the consent of the person concerned or some other legitimate basis laid down by law. Everyone has the right of access to data which has been collected concerning him or her, and the right to have it certified”. The processing must comply with the essential components of personal data protection, namely independent supervision and respect for the data subject’s rights.¹²

The right to personal data protection comes into play whenever personal data are processed. It is thus broader than the right to respect for private life, given that private life does not necessarily include all information on identified or identifiable persons. Any processing (i.e., collecting, recording, organising, structuring, storing, etc.) of personal data is subject to appropriate protections. Data protection concerns all kinds of personal data and data processing, irrespective of the relationship and impact on privacy. Processing of personal data may also infringe on the right to private life. However, it is not necessary to demonstrate an infringement on private life for data protection rules to be triggered.¹³

2.2. Data Protection under the GDPR

The territorial, material, and personal scope of the right to personal data protection is defined through secondary legislation¹⁴, namely the GDPR. Contrary to the ECHR and the CFR which create obligations for States, the GDPR directly imposes obligations on private parties. It became applicable on the 25th of May 2018 in all EU Member States. The GDPR applies to the “*processing of personal data wholly or partly by automated means and to the processing other than by automated means of personal data which form part of a filing system or are intended to form part of a filing system*”.¹⁵

The GDPR covers personal data, that is, any data related to an identified or an identifiable natural person. It also provides a stricter regime for data that is listed as “special categories” that may be

¹⁰ Article 16 TFEU provides the same wording.

¹¹ In that respect, while the case law of the European Court of Human Rights under Article 8 ECHR is obviously relevant when applying to the right of data protection, is not necessarily conclusive when assessing whether a situation is compliant with that Charter right. See “[t]he EU Charter of Fundamental Rights, A commentary”, Edited by Steve Peers, Tamara Harvey, Jeff Kenner and Angela Ward, Hart Publishing, 2014, page 235.

¹² Hustinx, P., EDPS Speeches & Articles, EU Data Protection Law: The Review of Directive 95/46/EC and the Proposed General Data Protection Regulation, July 2013.

¹³ Handbook on European data protection law, 2018 edition.

¹⁴ For the definition of secondary legislation see *supra*, footnote 2.

¹⁵ GDPR, art. 2(1).

regarded as “sensitive”. Given the activities envisaged under MobiDataLab, the GDPR will *prima facie* be applicable. However, if data have been anonymised, GDPR is not applicable. But it needs to be ensured that data cannot be returned to a “normal” state when relying on the way that personal details have been obfuscated.¹⁶ This means that the threshold for anonymisation is quite high.¹⁷

The key notions determining the material scope of application of GDPR will be clarified below.

2.2.1. Definitions

Personal data

Personal data is broadly defined as any type of information that relates to an identified or identifiable natural person (“data subject”).¹⁸ The CJEU has clarified that personal data is not limited to sensitive or private information, but potentially encompasses all types of information, both subjective and objective provided that it relates to the data subject.¹⁹ The ECHR has also interpreted the term “personal data” as not being limited to matters of the private sphere of an individual.²⁰

The information on a person can be clear, *i.e.*, directly identifying an individual (e.g., name, surname), or it can indirectly allow for the individual to be identified (e.g., by combining information on the specific hour a ticket is validated and footage from surveillance cameras). To determine whether a person is identifiable, a controller or another person must consider all reasonable means that are likely to be used directly or indirectly to identify the individual, such as, singling out. To determine whether identification is possible, one should consider all means that are reasonably likely to be used by the data controller (see below on the definition) or another person.²¹

When assessing such means, one should consider all factors at stake, such as the cost of conducting identification, the available technology, the risk of organisational dysfunctions (e.g., breaches of confidentiality duties), technical failures, etc.²² Moreover, the possibility of identification has to be assessed taking into account technological developments during the period for which the data will

¹⁶ GDPR, How to achieve and maintain compliance, Andrew Denley, Mark Foulsham, Brian Hitchen, Section 1.

¹⁷ For further details on anonymisation, please see section 4.2.3.

¹⁸ GDPR, art. 4(1).

¹⁹ CJEU, Case C-434/16, *Peter Nowak v. Data Protection Commissioner*, para. 34, ECLI:EU:C:2017:994.

²⁰ Case *Amann v. Switzerland*, Application no. 27798/95, para. 65.

²¹ GDPR, recital. 26.

²² *Ibid*; Article 29 Data Protection Working Party (2007) Opinion 4/2007 on the concept of personal data (01248/07/EN WP 136), p. 15.

be processed, keeping in mind that identification that may not be possible today given the current state of technology, may be possible in the future.²³

Processing

Processing of personal data means “any operation or set of operations which is performed on personal data or sets of personal data, whether or not by automated means”.²⁴ The concept of processing activities is very broad. Examples of processing activities include collection, recording, organisation, structuring, storage, adaptation or alteration, retrieval, consultation, use, disclosure, dissemination, alignment, restriction, erasure, or destruction of data.²⁵

Automated data processing concerns operations performed on “personal data wholly or partly by automated means”.²⁶ Practically, this means that any personal data processing through automated means with the help of, for example, a computer or a mobile device is covered under the GDPR. But processing is not restricted to automation only; processing personal data in a manual filing system, that is, a specially structured paper file also falls within the scope of the Regulation.²⁷

Data subject

The person whose personal data is protected under the GDPR is a living natural person, which is defined as the data subject. Legal persons do not benefit from legal protection under the GDPR.²⁸ The GDPR grants data subjects several rights: the right of access, the right to rectification, the right to erasure, the right to restrict processing, the right to data portability, the right to object, and the right not to be subject to a decision based solely on automated processing.²⁹

Controller

The controller is “any natural or legal person, public authority agency or other body which, alone or jointly with others, determines the purposes and means of the processing of personal data. Where the purposes and means of processing are determined by national or Community laws or regulations,

²³ *Ibid.*

²⁴ GDPR, art. 4(2).

²⁵ *Ibid.*

²⁶ GDPR, art. 2(1) and 4(2).

²⁷ *Supra*, ft.13, pp. 99-100.

²⁸ There have been however cases where legal entities were able to rely on Article 8 ECHR, if they are directly affected by a measure which breaches their right to respect for their “correspondence” or “home”. See Guide to the case-law of the ECtHR, Data protection, 31 December 2020, p. 8.

²⁹ GDPR, Article 15 onwards.

the controller or the specific criteria for his nomination may be designed by national or Community law”.³⁰

The assessment of controllership should be made based on the facts of a particular case. The key criterion to assess who is a controller is to designate the person who determines the “purposes” (the “why”) and the “means” (the “how”) of the processing of personal data.³¹ It seems however the purpose may take precedence over the means. As such, determining the purpose of the processing, in any case, leads to a qualification as a controller. Determining the means would lead to control only when it concerns the essential means, such as which data is processed, the duration of the processing, which third parties have access to the data.³² The determination of the technical and organisational elements of the means (e.g., which hardware or software to use) does not necessarily imply control and can hence be done exclusively by the processor.³³

Under the GDPR, both natural, legal persons and a public authority can be considered as a controller. But normally it would be the company or a body that would qualify as a controller, rather than a specific individual within the company or the body.³⁴

Joint controllership

When two or more parties jointly determine the purpose and means of processing, they are considered joint controllers.³⁵ “Jointly” must be interpreted as meaning “together with” or “not alone”, in different forms and combinations.³⁶ Joint controllership exists when the parties decide together to process data for the same or common purpose. Joint controllership also requires that two or more entities have exerted influence over the means of the processing. This does not mean that, for joint controllership to exist, each entity involved needs to determine all the means in each case. Indeed, as clarified by the CJEU, different entities may be involved at different stages of that processing and to different degrees.³⁷

³⁰ GDPR, art. 4(7).

³¹ European Data Protection Board, Guidelines 7/2020 on the concepts of controller and processor in the GDPR, Version 1.0, section 2.1.4, p. 13.

³² *Ibid.* p.14.

³³ *Ibid.*

³⁴ *Ibid.*, section 2.1.1, p.9.

³⁵ GDPR, Articles 4(7) and 26.

³⁶ *Supra*, ft. 31, section 3.2.1, p. 17.

³⁷ *Ibid.*, section 3.2.2.2, p. 19.

Processor

The processor is a natural or legal person, public authority, agency, or other body which processes personal data on behalf of the controller.³⁸ In principle, there is no limitation as to which type of actor might assume the role of a processor. It might be an organisation, but it might also be an individual. Two characteristics define who can qualify as a processor: i) being a separate entity in relation to the controller, and ii) processing personal data on the controller's behalf.³⁹ However, it has been noted that not every service provider that processes personal data while delivering a service is a "processor" within the meaning of the GDPR. The role of a processor is granted not from the mere processing of data but its concrete activities in a specific context. It is the nature of the service that will determine whether the processing activity amounts to the processing of personal data on behalf of the controller within the meaning of the GDPR.⁴⁰

³⁸ GDPR, art. 4(8).

³⁹ *Supra*, ft. 31, section 4, p. 24.

⁴⁰ *Ibid.*

3. Sources and Uses of Location Data

This document focuses on regulations, privacy risks, and protection techniques that apply to those mobility data which contain personal information. However, we cannot ignore the protection of economic interests related to operational information collected, processed, and/or shared by private organisations, which may additionally be subject to industrial property rights.

The applicability of the protection mechanisms described in this document depends on the activities to be performed on location-based mobility data, and the sources from which it is collected. We distinguish between location-based services (LBS), where service providers can collect location information directly from their users for primary (e.g., route planning) and secondary (e.g., modelling of traffic patterns) uses, for which data is stored and processed. Trajectory microdata sharing, coming from data collected through LBS or other means, and useful for processing and data mining. And aggregated mobility data, which can be used to communicate patterns and results or as side information for further analyses.

Deliverable D2.9 (Use Cases Definition V1) describes the initially considered use cases to demonstrate the validity and utility of the MobiDataLab proposed platform. These use cases are classified as operational use cases and research use cases. Next, we enumerate these use cases, including their points in common and their particularities so that we can provide further recommendations towards the design of the MobiDataLab architecture and to prioritise the development of protection mechanisms in terms of their applicability to the considered use cases. Full details on these use cases are available in D2.9.

Use cases for operations include:

- Optimisation of transport flow and estimated time of arrival;
- Emission reporting;
- Analytics and learning;
- Re-use of transport data for journey planners;
- Mobility as a service.

Use cases for research include:

- Geodata sharing applied to transport (for example, inclusive and environmentally friendly transport);
- Transport data sharing within the Linked Open Data Cloud (to link, for example, mobility and tourism data).

D2.9 additionally lists the data required to implement each of these use cases. We next provide a consolidated list of these data and describe the possible sources from which to obtain them, which in turn influence the sensitivity of the data and on the applicable protection mechanisms that could be employed to reduce the risks of disclosure.

Table 1: Location data required in the use cases. Includes whether the data are personal information or not, and which Section(s) in this document deal(s) with the protection of such data.

Type of data	Description and sourcing	Personal	Related Section
Driver location data	Driver location data in real time for route planning, ETA, emission reporting and other location-based services. These data are directly sourced from the users, via their mobile phones or GPS-enabled vehicles.	YES	5
GPS traces	GPS traces, also known as trajectory microdata in this document, consist of databases in which each record corresponds to an individual and contains the location data from this individual in each period of time. These traces are obtained by collecting and merging user-provided data.	YES	6
Static map data or transport network topology	Static map data or transport network topology, corresponding to maps and roads. These data are offered by public entities such as municipalities and departments of transport. Organisations such as Google and OpenStreetMap provide access to this kind of data	NO	N/A
Real-time traffic data	Real-time traffic data, including disruptions and alerts and planned events. This information is also typically provided by public authorities in real time, although it is common for service providers to offer real-time or aggregated traffic information obtained through crowdsourcing. Service providers that continuously monitor the location of their users (e.g., Google via Android phones), can obtain traffic information aggregating the speed, location, and density of their users, and enriching this data with publicly available information.	YES/NO	5, 6, 7
Points of interest	Points of interest can be obtained from public authorities, such as departments of tourism, by companies announcing themselves or by crowdsourcing. Pols can also be obtained by analysing highly visited locations in trajectory datasets.	YES/NO	6, 7
Vehicle data	Vehicle data, fuel type and load are useful information for emissions reporting. These data must be supplied by users or companies.	YES/NO	4
Public transport data	Public transport data is offered by public transport operators, although they can also be obtained through crowdsourcing from public transport users, in real time or not.	YES/NO	5, 6, 7

Information collected directly from users, such as real-time location data or (Global Positioning System) GPS data are highly sensitive data which requires special attention. These data can be

either protected locally by the users before sending them to the service provider, or by the controller if any sharing or release is intended. We enumerate several strategies for the protection of these kinds of data in Sections 5 and 6, where we deal with location-based services and trajectory microdata. Other kinds of data, such as real-time traffic or public transport data are also potentially sensitive when they are obtained via crowdsourcing. Section 7 describes some of the risks of publishing or sharing aggregate data, such as heatmaps. Data such as maps are considered public information and do not require any additional protection.

4. General Data Privacy Measures

According to the GDPR, controllers⁴¹ need to address data protection issues already at the design phase of their activities and business practices to meet the requirements of the Regulation and protect the rights of data subjects⁴² (“*privacy by design*”). To achieve privacy by design, controllers also need to ensure that they abide by the data protection principles: lawfulness, fairness, and transparency of processing; purpose limitation; data minimisation; accuracy; storage limitation; integrity and confidentiality (security), and accountability.⁴³

Controllers are required to deploy *appropriate* technical and organisational measures and *necessary* safeguards both at the time of processing and when determining the means for processing. In implementing these measures, the controller needs to consider the state of the art, the costs of implementation, the nature, scope, and purposes of personal data processing, and the risks and severity for the rights and freedoms of the data subject.⁴⁴

Examples of technical or organisational measures and safeguards include pseudonymisation of personal data (e.g. by using encryption with a secret key, hash function); storing personal data available in a structured, commonly machine-readable format; enabling data subjects to intervene in the processing; providing information about the storage of personal data; having malware detection systems; training employees about basic “cyber hygiene”; establishing privacy and information security management systems, obligating processors contractually to implement specific data minimisation practices, etc.⁴⁵ Standards, best practices, and codes of conduct that are recognised by associations and other bodies representing categories of controllers can help determine appropriate measures. However, the controller must verify the appropriateness of the measures for the processing in question.⁴⁶

⁴¹ For the definition of a “controller”, see section 2.2.1.

⁴² For an overview of those rights, see section 2.1.

⁴³ GDPR, Article 5 and recital 39.

⁴⁴ GDPR, Article 25 (1) and recital 78. The GDPR also requires “data protection by default”, meaning that companies/organisations should ensure that personal data is processed with the highest privacy protection (for example only the data necessary should be processed, short storage period, limited accessibility) so that by default personal data isn’t made accessible to an indefinite number of persons. For further information see: https://ec.europa.eu/info/law/law-topic/data-protection/reform/rules-business-and-organisations/obligations/what-does-data-protection-design-and-default-mean_en.

⁴⁵ EDPB, Guidelines 4/2019 on Article 25 Data Protection by Design and by Default Version 2.0 Adopted on 20 October 2020, p.6.

⁴⁶ *Ibid.*

In this Section, we introduce data protection design strategies, measures, tools, and techniques applicable to general scenarios and including all manners of data, not restricted to mobility or transportation data.

4.1. Data protection by design

4.1.1. Design strategies

The following design strategies aid service providers in complying with the GDPR. Thus, when designing a software product which deals with personal information, system designers and service providers (which act as data controllers and possibly as data processors) should take the following principles into account. These design principles reflect the data protection principles enshrined in the GDPR and analysed above, were first described by Hoepman⁴⁷ and are divided in two groups: the first four being data-oriented principles and the last four, process-oriented or operational principles.

The **Minimise** design strategy consists of collecting no more data than those which are strictly necessary to the correct operation of the service that is being provided, along with information otherwise required by state legislation or authorities (e.g., by fiscal authorities).

The **Hide** design strategy refers to the use of encryption, de-identification, anonymisation and/or pseudonymisation technologies wherever possible to ensure the confidentiality of data, both in transit and in storage.

The **Separate** strategy states that, if possible, personal information should be processed in a distributed way. This strategy can be translated as both ensuring that databases including information about the same individual, but with different data (e.g., shopping history and billing information), are not linked within the organisation; and, as a more extensive interpretation, recurring to decentralised or peer-to-peer architectures when possible, with clients processing their data locally and sharing information with the rest of the clients or the service provider in a need-to-know basis.

The **Aggregate** design strategy indicates that personal information should be processed at the highest possible level of aggregation, resorting to high granularity (or individual) information only if mandatory.

The **Inform** principle relates to transparency. Respondents should be informed of all collection, processing, and possible sharing of their personal data, and of the purpose for such activities.

⁴⁷ Hoepman, J.H., 2014, June. Privacy design strategies. In IFIP International Information Security Conference (pp. 446-459). Springer, Berlin, Heidelberg.

The **Control** principle states that respondents should retain control of their data and be allowed to consult, modify, and delete any information about them held by an organisation.

The **Enforce** principle tells organisations to have privacy policies in place and enough personnel and organisational structures and effort to ensure the policies are enforced.

The **Demonstrate** principle relates to accountability. Organisations ought to document all collection and processing of personal information, along with all measures put in place (e.g., following the strategies above) to ensure those data are kept secure and that the respondent's privacy is preserved.

4.2. Tools and Techniques

In this Section, we introduce a series of tools and techniques to ensure the protection of personal data. We include both cryptographic and non-cryptographic mechanisms.

Cryptographic mechanisms may be used for data pseudonymisation, *i.e.*, replacing one attribute in a record – which makes it possible to identify the data subject – by another, and keeping those attributes separate, under technical or organisational measures.

Article 25 of the GDPR which addresses data protection by design, as mentioned above, explicitly refers to pseudonymisation as an example of an appropriate technical and organisational measure that controllers should implement to accommodate the data protection principles and integrate the necessary safeguards.⁴⁸

The process of pseudonymisation must not be confused with the process of anonymisation, where all links to identifying the person are broken.

4.2.1. Cryptographic tools

4.2.1.1. Encryption

In general terms, data encryption refers to some operations performed on data which render them unintelligible to everyone except for those entities with the right to access to the data. This access right normally refers to the entity having the decryption key. The way decryption keys work, are managed, and are distributed lets us classify encryption schemes in two categories: symmetric or secret-key encryption and asymmetric or public-key encryption.

⁴⁸ Handbook on European data protection law, 2018 edition, pp. 131-132.

Symmetric encryption

Symmetric or secret-key encryption schemes take a piece of data (a *message* or *plaintext*) and a *key* of a given length and return a *ciphertext*, which can only be decrypted by applying the decryption operation using the same *key* that was used to encrypt the *plaintext*. Thus, to communicate a secret message between two entities, these entities need to share a *secret key*.

Formally, a symmetric encryption scheme includes a probabilistic encryption algorithm $Enc: \mathcal{M} \times \mathcal{K} \rightarrow \mathcal{C}$, which takes a message m from a message space \mathcal{M} and a key k from a key space \mathcal{K} and returns a ciphertext c ; and a deterministic decryption algorithm $Dec: \mathcal{C} \times \mathcal{K} \rightarrow \mathcal{M}$, which, given a ciphertext c and a key k returns a message m . For correctness, given some message m and an honestly generated key k , $Dec(Enc(m, k), k) = m$, that is, the decryption of an honestly encrypted message will return the same message. Security notions for symmetric encryption schemes refer to *indistinguishability* against attacks from adversaries with different capabilities and can be summarised as follows: an attacker that does not know the secret key (and with specific capabilities) should not be able to distinguish a ciphertext from a random string⁴⁹. Symmetric encryption schemes are built from information-theoretic principles, and thus are considered secure against quantum computers.

Symmetric encryption schemes are further divided in two families: block ciphers, which are used to encrypt information in blocks of fixed size, and stream ciphers, which are used to encrypt data streams on bit at a time. Popular block cipher schemes include DES, 3DES, IDEA, Blowfish, and Rijndael/AES, AES-256 being the current standard block cipher recommended for modern applications. For stream ciphers, ChaCha20 is the current recommended scheme.

Symmetric encryption schemes are very efficient and are usually implemented in hardware.

Public-key encryption

Asymmetric or public-key encryption schemes ease key management by removing the requirement for both sender and receiver to share the same key ahead of time. Instead, each entity is associated with two keys: a public key used by an encryptor; and a private key used to decrypt messages encrypted under the corresponding public key. Thus, to send a message it is no longer necessary to know the decryption key. Additionally, the public key may be transmitted or published in the clear so there is no requirement for secure channels before transmitting a message.

Formally, a public-key encryption scheme consists of three algorithms. A probabilistic key generation algorithm $KeyGen: \mathbb{N}^+ \rightarrow \mathcal{PK} \times \mathcal{SK}$, that given a security parameter (e.g., the key length), returns a key pair consisting of a public key pk and a secret key sk ; a (probabilistic) encryption algorithm $Enc: \mathcal{M} \times \mathcal{PK} \rightarrow \mathcal{C}$, that given a message m and a public key pk returns a ciphertext c ; and a deterministic decryption algorithm $Dec: \mathcal{C} \times \mathcal{SK} \rightarrow \mathcal{M}$, that given a ciphertext c and a secret key sk returns a message m . For correctness, given an honestly generated key pair $(pk, sk) \leftarrow KeyGen(\lambda)$, for security parameter λ , $Dec(Enc(m, pk), sk) = m$. Similar security notions of *indistinguishability* are

⁴⁹ Bellare, M., Desai, A., Joripii, E. and Rogaway, P., 1997, October. A concrete security treatment of symmetric encryption. In Proceedings 38th Annual Symposium on Foundations of Computer Science (pp. 394-403). IEEE.

defined for public-key encryption. Current public-key encryption algorithms are based on number-theoretic problems (mainly the discrete logarithm problem and the integer factorisation problem), which are known to be weak against quantum computers. National Institute of Standards and Technology (NIST) is currently holding an open challenge to find new standards for quantum-resistant public-key encryption⁵⁰.

The current public-key encryption standard is RSAES-OAEP, using the RSA cryptosystem⁵¹, based on the integer factorisation problem, together with the OAEP padding scheme⁵², which provides RSA with semantic security (*i.e.*, randomises the encryption algorithm). Recommended key lengths for RSA stand currently at 3,092 bits.

Public-key encryption algorithms are significantly less efficient than symmetric encryption schemes.

Key exchange protocols and hybrid cryptosystems

As discussed above, symmetric key encryption schemes are highly efficient, but require every pair of entities to share an encryption-decryption key, which must be transmitted using a secure channel prior to any exchange. Thus, managing keys becomes a very complex matter. On the other hand, public-key encryption solves this complexity, but at the cost of inefficient encryption and decryption procedures. Therefore, using public-key encryption schemes to send big volumes of data is highly impractical. Key exchange protocols and hybrid cryptosystems allow us to get the best of both worlds.

The Diffie-Hellman key exchange⁵³ (DHKE) protocol allows two entities to agree on a shared key using an insecure communication channel without the need to exchange any prior information. The DHKE protocol is based on the discrete logarithm problem. This protocol is currently the standard key exchange protocol as defined in PKCS#3 and serves as the building block for any secure internet communication, being part of TLS, HTTPS, and end-to-end encryption schemes. Being based on the discrete logarithm problem, the DHKE protocol is not resistant to quantum cryptanalysis, and new standards for quantum-resistant key exchange protocols are currently being sought by NIST.

An alternative method to benefit from the both symmetric and public-key cryptosystems are key encapsulation mechanisms, also known as hybrid cryptosystems. In hybrid cryptosystems, a public-key encryption scheme is used to encrypt a secret key belonging to a symmetric key cryptosystem, which is used to encrypt the message to be transmitted. For example, if we want to transmit message m , we generate a random key k and encrypt the message using AES with the key k . This key is then

⁵⁰ Post-quantum cryptography standardisation – <https://csrc.nist.gov/Projects/post-quantum-cryptography/round-3-submissions>

⁵¹ Rivest, R.L., Shamir, A. and Adleman, L., 1978. A method for obtaining digital signatures and public-key cryptosystems. Communications of the ACM, 21(2), pp.120-126.

⁵² Bellare, M. and Rogaway, P., 1994, May. Optimal asymmetric encryption. In Workshop on the Theory and Application of Cryptographic Techniques (pp. 92-111). Springer, Berlin, Heidelberg.

⁵³ Diffie, W. and Hellman, M., 1976. New directions in cryptography. IEEE transactions on Information Theory, 22(6), pp.644-654.

encrypted using RSA with the public key of the desired recipient. We then send the encryption of m under key k and the encryption of k under the public key of the recipient. The recipient can use her private key to recover k and use it to recover the message m . These hybrid cryptosystems are typically used for secure e-mail.

4.2.1.2. Computations on encrypted data

While encryption is the default (recommended) option to keep data at rest and in transit confidential, we often need to perform computations on these data. However, most data analyses are incompatible with most encryption procedures: users typically require data in clear form to analyse them.

A possible candidate to allow computations on data while keeping the information private is anonymisation, but some cryptographic techniques also allow us to perform (some limited) computations on encrypted data and are usually part of larger systems, such as privacy-preserving data mining.

Functional encryption

Functional encryption⁵⁴ extends public-key encryption to allow the holder of a secret key to learn some specific function of an encrypted message, but nothing else. In some cases, this function could return the message itself (e.g., as per standard public key encryption, or return the message only if some additional criteria are met) or instead may produce the output of some computation specified by the message and secret key. Functional encryption is very powerful, but definitionally challenging in the general case since the functionality itself can be varied.

(Partially) Homomorphic Encryption

Some encryption schemes are homomorphic in nature. Given two ciphertexts encrypting two plaintexts, certain operations can be performed on the ciphertexts such that the result can be decrypted to produce the outcome of applying an operation (not necessarily the same) on the plaintexts themselves. Thus, some computations can be performed on encrypted data. Schemes that exhibit homomorphic properties for a specific operation are known as partially homomorphic encryption schemes. Examples of this class are ElGamal⁵⁵ and Paillier⁵⁶.

⁵⁴ Boneh, D., Sahai, A. and Waters, B., 2011, March. Functional encryption: Definitions and challenges. In Theory of Cryptography Conference (pp. 253-273). Springer, Berlin, Heidelberg.

⁵⁵ ElGamal, T., 1985. A public key cryptosystem and a signature scheme based on discrete logarithms. IEEE transactions on information theory, 31(4), pp.469-472.

⁵⁶ Paillier, P., 1999, May. Public-key cryptosystems based on composite degree residuosity classes. In International conference on the theory and applications of cryptographic techniques (pp. 223-238). Springer, Berlin, Heidelberg.

On the other hand, if the set of permissible operations enable arbitrary computations to be performed, then the schemes are referred to as fully homomorphic^{57,58}. Although fully homomorphic schemes are in principle very powerful, currently available instances also involve very substantial overhead and storage expansion.

4.2.1.3. Multiparty Computation

Secure multiparty computation protocols allow a set of parties to compute a joint function of their inputs in a secure way without requiring a trusted third party. During the execution of the protocol, the parties do not learn anything about each other's input except what is implied by the output itself.

A general solution for the secure computation of functions among two players was introduced by Yao⁵⁹. The main idea of these protocols was to describe the function as a circuit, and to compute every gate of the circuit in a secure way. This idea was extended to the multi-partite setting by Goldreich *et al.*⁶⁰ They showed how to create a secure multiparty computation protocol that allows playing any game and does not leak any information if most of the participants are honest. The first unconditionally secure multi-party computation protocols were presented by Ben-Or *et al.*⁶¹ and Chaum *et al.*⁶² These authors gave protocols to compute any arithmetic function in a secure way when at least two thirds of the parties are honest.

⁵⁷ Gentry, C., 2009, May. Fully homomorphic encryption using ideal lattices. In Proceedings of the forty-first annual ACM symposium on Theory of computing (pp. 169-178).

⁵⁸ Gentry, C., Sahai, A. and Waters, B., 2013, August. Homomorphic encryption from learning with errors: Conceptually-simpler, asymptotically-faster, attribute-based. In Annual Cryptology Conference (pp. 75-92). Springer, Berlin, Heidelberg.

⁵⁹ Yao, A.C.C., 1986, October. How to generate and exchange secrets. In 27th Annual Symposium on Foundations of Computer Science (sfcs 1986) (pp. 162-167). IEEE.

⁶⁰ Goldreich, O., 1987, January. Towards a theory of software protection and simulation by oblivious RAMs. In Proceedings of the nineteenth annual ACM symposium on Theory of computing (pp. 182-194).

⁶¹ Ben-Or, M., Goldwasser, S. and Wigderson, A., 1991. Completeness theorems for non-cryptographic fault-tolerant distributed computation. In Providing Sound Foundations for Cryptography: On the Work of Shafi Goldwasser and Silvio Micali (pp. 351-371).

⁶² Chaum, D., Crépeau, C. and Damgard, I., 1988, January. Multiparty unconditionally secure protocols. In Proceedings of the twentieth annual ACM symposium on Theory of computing (pp. 11-19).

4.2.1.4. Integrity and authenticity

Identity, authentication, and access control are central components of secure systems. It is important that data assets be accessible only to authorised parties. On the one hand, a sound authentication and authorisation infrastructure prevent data breaches. On the other hand, it allows responsibilities to be attributed in case of a breach, which contributes to a transparent data processing environment.

Several methods exist to verify the identity of individuals, that is, to authenticate them. Some of them allow for the authentication of users without disclosing their identity.

Hash functions

Hash functions⁶³ are compression functions that take an arbitrary length string and output a shorter string. For cryptographic applications we require additional security properties to ensure that the mapping from long inputs to short outputs does not produce too many collisions, and that the mapping can be hard to work out. The required security properties are as follows:

- Collision Resistance: it should be computationally infeasible to find two distinct messages that hash to the same value.
- Pre-image Resistance: it should be computationally infeasible to find a message such that its hash is a desired value, in other words, hash functions are one-way functions (computationally infeasible to invert).
- Second Pre-image Resistance: given a value and its hash result, it should be computationally infeasible to find another value that hashes to the same value.

There are several standardised hash functions that achieve different levels of security in the above notions. The current NIST standards are SHA-256/512 and SHA-3.

Digital Signatures

In paper documents, handwritten signatures guarantee the authenticity of the document, and the signer cannot repudiate it. Moreover, the paper support gives some protection against manipulation: deletions and additions can be detected, at least by an expert. Digital signatures were created to guarantee the authenticity and integrity in the case of electronic communications, and to avoid their repudiation. Digital signatures were made possible by the deployment of public-key encryption. In addition, digital signatures and the public key infrastructure can be used to provide authentication of individuals.

Formally, a digital signature scheme consists of three algorithms. A probabilistic key generation algorithm $KeyGen: \mathbb{N}^+ \rightarrow \mathcal{VK} \times \mathcal{SK}$, that given a security parameter (e.g., the key length), returns a key pair consisting of a verification key vk and a signing key sk ; a (probabilistic) signature algorithm $Sign: \mathcal{M} \times \mathcal{SK} \rightarrow \mathcal{S}$, that given a message m and a signing key sk returns a signature s ; and a deterministic verification algorithm $Ver: \mathcal{M} \times \mathcal{S} \times \mathcal{VK} \rightarrow \{0,1\}$, that given a message m , a signature s and a verification key vk indicates whether the signature is correct (1) or not (0). For correctness, given an honestly generated key pair $(vk, sk) \leftarrow KeyGen(\lambda)$, for security parameter λ ,

⁶³ Katz, J. and Lindell, Y., 2020. Introduction to modern cryptography. CRC press.

$Ver(m, Sign(m, sk), vk) = 1$. Note that, typically, a hash function is applied to the messages to reduce their size. Security notions for digital signatures relate to *unforgeability*, meaning an attacker without knowledge of the signing key cannot produce valid signatures. Digital signature schemes are based on the same or similar number-theoretic problems as public-key encryption and are therefore also not secure against quantum attacks. NIST is also currently holding an open challenge to find new standards for quantum-resistant signature schemes⁶⁴.

Current signature standards include DSA⁶⁵, based on the ElGamal public-key encryption scheme and ECDSA⁶⁶, which is the same algorithm but based on elliptic curve groups instead of integer groups.

Next, we enumerate specific classes of digital signatures that enable authentication while being compatible with some user anonymity.

- Blind signatures⁶⁷ are considered particularly useful for electronic payment systems, electronic voting schemes and token-based access control mechanisms; a user may obtain a signature (e.g., a signed coin from a bank) such that the signer does not know the contents of the message and cannot produce further valid signatures.
- In a group signature scheme⁶⁸, a set of users, called members of the group, can issue signatures of arbitrary messages on behalf of the group. A verifier can check the validity of the signature using the group public key. The main interest in this kind of signature is that it ensures the privacy of signers against potential verifiers, because a potential verifier cannot distinguish two signers from the same group.
- A requirement of group signatures is the support for membership revocation of misbehaving members without the need to update the group public key. To facilitate member revocation, some members, called group managers, are endowed with the capability to revoke membership.
- Identity-based signature schemes, theorised by Shamir in 1984⁶⁹ and with the first concrete protocol, based on the Weil pairing, shown by Boneh and Franklin in 2001⁷⁰, allow public keys

⁶⁴ Post-quantum cryptography standardization – <https://csrc.nist.gov/Projects/post-quantum-cryptography/round-3-submissions>

⁶⁵ Kerry, C.F., and Gallagher, P.D., 2013. Digital signature standard (DSS). FIPS PUB, pp.186-4.

⁶⁶ Johnson, D., Menezes, A. and Vanstone, S., 2001. The elliptic curve digital signature algorithm (ECDSA). International journal of information security, 1(1), pp.36-63.

⁶⁷ Chaum, D., 1983. Blind signatures for untraceable payments. In Advances in cryptology (pp. 199-203). Springer, Boston, MA.

⁶⁸ Chaum, D. and Van Heyst, E., 1991, April. Group signatures. In Workshop on the Theory and Application of Cryptographic Techniques (pp. 257-265). Springer, Berlin, Heidelberg.

⁶⁹ Shamir, A., 1984, August. Identity-based cryptosystems and signature schemes. In Workshop on the theory and application of cryptographic techniques (pp. 47-53). Springer, Berlin, Heidelberg.

⁷⁰ Boneh, D. and Franklin, M., 2001, August. Identity-based encryption from the Weil pairing. In Annual international cryptology conference (pp. 213-229). Springer, Berlin, Heidelberg.

to be arbitrary strings of some length, called identities. These strings are associated with a user and reflect some aspect of her identity, e.g., her email address. The corresponding secret key is then computed by a trusted entity taking as input the user's identity and, possibly, some other secret information, and is sent to the user through some secure channel. Identity-based public key signature schemes offer considerable flexibility in key generation and management.

- Attribute-based signatures generalise identity-based signatures in that, instead of having the users' identities as credentials, they use properties, or attributes, of the users as the latter's' credentials (in the attribute-based setting, the identity is one more attribute of the user). Attribute-based signatures were introduced by Shanqing and Yingpei in 2008⁷¹. In attribute-based signatures (and encryption) schemes, the users receive private key shares associated with their credentials, such as their name, age, country of residence, having or not a driving licence, place of work, etc. Digital signatures are produced with respect to some function of the users' credentials, typically called a policy.

Zero-Knowledge Proofs

Zero-knowledge proofs⁷² allow a prover to convince a verifying party of the truth of a statement without revealing any information other than the truth of the statement. In particular, if the statement requires the prover to hold some secret information, then the verifier does not learn this information—it is possible to prove knowledge of a secret without revealing the secret itself. Statements that only prove possession of a secret are known as zero-knowledge proofs of knowledge. Proofs can be either interactive or non-interactive depending on whether the parties can communicate during the proof. In general, non-interactive zero knowledge (NIZK) proofs⁷³ are considered more difficult since they cannot use interactive challenge-response protocols. zkSNARKs and zkSTARKS, used in several cryptocurrencies, are in this family of proofs. Whereas zero-knowledge proofs can be rather inefficient, non-interactive proof systems built on bilinear groups⁷⁴ are particularly efficient for group-dependent problems where the secrets are group elements or the exponents of a group element. As many useful cryptographic schemes are built using bilinear pairings, particularly functional encryption, such a proof system can be very useful for proving knowledge of a cryptographic secret without revealing it.

⁷¹ Shanqing, G. and Yingpei, Z., 2008, April. Attribute-based signature scheme. In 2008 International Conference on Information Security and Assurance (ISA 2008) (pp. 509-511). IEEE.

⁷² Ben-Or, M., Goldreich, O., Goldwasser, S., Håstad, J., Kilian, J., Micali, S. and Rogaway, P., 1988, August. Everything provable is provable in zero-knowledge. In Conference on the Theory and Application of Cryptography (pp. 37-56). Springer, New York, NY.

⁷³ Blum, M., Feldman, P. and Micali, S., 2019. Non-interactive zero-knowledge and its applications. In Providing Sound Foundations for Cryptography: On the Work of Shafi Goldwasser and Silvio Micali (pp. 329-349).

⁷⁴ Groth, J. and Sahai, A., 2008, April. Efficient non-interactive proof systems for bilinear groups. In Annual International Conference on the Theory and Applications of Cryptographic Techniques (pp. 415-432). Springer, Berlin, Heidelberg.

Zero-knowledge proofs can be used to authenticate users holding cryptographic devices, such as smartcards, without leaking any information about these users except that they hold a valid card.

4.2.2. Private Communications

This section discusses the protection of communication channels. First, it describes end-to-end encryption, which provides confidentiality of communications. It then introduces anonymous channels. Having discussed mechanisms that allow users to be authenticated without revealing their identities, it is logical to discuss communication channels that do not reveal their address, which is also part of their identity.

4.2.2.1. End-to-End Encryption

End-to-end encryption refers to the encryption of messages exchanged by two or more parties without the intervention of a centralised server. The centralised server may exist and support the transport of the messages, but all this server sees is encrypted content. This behaviour is the opposite of the traditional message exchange protocols, in which the messages are only encrypted while in transit from the parties to the central server or from the central server to the parties.

End-to-end encryption is typically supported by having all participants have a key pair from a public-key encryption scheme. The centralised server, in addition to supporting the exchange of messages, works as a public-key repository, where users can find the public keys of the users to whom they want to send messages. Once a user has obtained another user's public key, she can use this public key to encrypt the messages, which will only be decryptable by the owner of the corresponding private key. A more efficient variant is for users to exchange random session keys for symmetric encryption by enciphering them under their public-private pairs and then encrypting the messages with a symmetric encryption scheme under these random temporal session keys.

4.2.2.2. Anonymous Channels

Anonymous channels allow users to hide their address (e.g., the Internet Protocol (IP) address) to the service provider they are communicating with. Examples of anonymous channels include mixnets and onion routing.

A mix network or mixnet is a routing protocol in which each of the network nodes shuffles (and re-encrypts) all received messages before sending them to the next node⁷⁵. The shuffling process is

⁷⁵ Chaum, D.L., 1981. Untraceable electronic mail, return addresses, and digital pseudonyms. Communications of the ACM, 24(2), pp.84-90.

kept secret by each mix server. Additionally, the sender of the message might successively encrypt the message with each of the mix servers' public keys. If that is the case, each mix server will have to decrypt each of the encryption layers (as if peeling an onion) until the destination of the message. The ToR network⁷⁶ is an example of this operation.

4.2.3. Anonymisation

On some occasions, we require some data processing operations that are not compatible with encryption, even with encryption schemes that allow some computations. Additionally, data collected by some organisations (governmental, medical, etc.) might be of interest to researchers, in which case these datasets can be released to research organisations or even made public. In all these cases, it is important that when such data contain information about individuals some measures are taken to prevent the release of sensitive information about such individuals. These measures are collectively known as anonymisation measures. Anonymous data is not personal data and hence falls outside the scope of the GDPR.

A dataset can be seen as a matrix with n rows that correspond to single individuals and k columns, which correspond to attributes about the individuals, attributes in a dataset are classified as follows:

- **Identifiers** are attributes that unambiguously identify individuals, for example, a passport number;
- **Quasi-identifiers** are attributes that, while not being identifiers themselves, can identify single persons when combined. Examples of this are date of birth, city of birth or, in location data, position of home and position of workplace;
- **Confidential attributes** are those which contain sensitive information about an individual, for example the salary, religion, or political affiliation;
- **Non-confidential attributes** are the rest of attributes about an individual that do not impact on their privacy.

Anonymisation measures aim to produce modified versions of the original datasets which are as close as possible to the original datasets but that do not allow attackers to identify individuals or to infer confidential attributes about concrete individuals.

⁷⁶ Dingledine, R., Mathewson, N. and Syverson, P., 2004. Tor: The second-generation onion router. Naval Research Lab Washington DC.

Table 2: *Example microdata table: The Name attribute is an Identifier, Age and Zipcode are quasi-identifiers, while the Diagnostic attribute is a Confidential attribute.*

Name	Age	Zipcode	Diagnostic
Alex	20	15k	Bronchitis
Bob	25	42k	Pneumonia
Jane	33	71k	Flu
Cathy	38	25k	Gastritis
Eva	44	56k	Emphysema
Frank	47	18k	Dyspepsia
David	53	31k	Bronchitis
Helen	61	35k	Flu

The threshold for anonymisation has been set high. According to the Article 29 Working Party⁷⁷, “an effective anonymisation solution should prevent all parties from singling out an individual in a dataset, from linking two records within a dataset (or between two separate datasets) and from inferring any information in such dataset. Therefore, removing directly identifying elements in itself is not enough to ensure that identification of the data subject is no longer possible.”⁷⁸ This means that if a controller de-identifies a dataset but at the same time keeps the original (identifiable) dataset – the de-identified dataset is unlikely to be considered truly anonymous but would still qualify as personal data.⁷⁹

The Article 29 Working Party assessed the strengths and weaknesses of several anonymisation (and other) techniques⁸⁰ by focusing on the following three risk factors:

- singling out, that is the “the possibility to isolate some or all records which identify an individual in the dataset”;
- linkability, that is the “ability to link at least two records concerning the same data subject or a group of data subjects (either in the same database or in two different databases). [...]”;

⁷⁷ This is the supervisory authority that was replaced by the Data Protection Supervisory Board and essentially ensures the same function (uniform interpretation of the GDPR).

⁷⁸ Article 29 Data Protection Working Party (2014). Opinion 05/2014 on anonymization techniques (829/14/EN WP216), p.9.

⁷⁹ *Ibid.*

⁸⁰ *Ibid*, section 3, p. 11.

- inference, which is “the possibility to deduce with significant probability the value of an attribute from the values of a set of other attributes”.

We describe these risks in detail in the next Section.

According to the Article 29 Working Party, an anonymisation technique that provides a solution against these three risks would be sufficiently robust against re-identification by the data controller or a third party considering the means they are reasonably likely to employ.⁸¹ The techniques assessed by the Article 29 Working Party and the outcome of such assessment are provided in Annex 1. Important to note is that the Article 29 Working Party concludes that each technique assessed fails (by itself) to address all three abovementioned risk factors.⁸² It does however point out that careful engineering may allow combining several techniques to enhance the robustness of the anonymisation outcome. The optimal solution should ultimately be decided on a case-by-case basis.⁸³

4.2.3.1. Risks and attacks

We consider attacks in which an attacker, with some background knowledge about some individual (e.g., knows some quasi-identifiers of the victim, such as their age or their place of birth) can identify such individual in a released dataset and, thus, learning some additional information about the individual. Note that the presence of the individual in a dataset might itself be a violation of the individual’s privacy. For example, if an individual is found in a dataset containing cancer treatments, one can infer that such individual is being treated for cancer, and thus that the individual has cancer. The risks here is that the identity of individuals or some attributes of such individuals are disclosed.

In record linkage attacks, the attacker knows some quasi-identifiers about some victim and checks for combinations of quasi-identifier attribute values in the released data base. If the attacker can find a small group of individuals with the same combination of quasi-identifier attribute values as her victim, then the attacker is able to link a record in the dataset to her victim with high confidence. k -Anonymity ensures that a minimum number of k individuals share the same quasi-identifier attribute values, and thus limits the risk of linkage to $1/k$.

Attribute linkage attacks are those in which an attacker learns some information on a confidential attribute from her target individual. Again, the attacker knows some quasi-identifier attribute values of her victim and proceeds in the same way as in the record linkage attacks. In case all individuals that share the same quasi-identifiers also share some confidential attribute, the attacker learns that confidential attribute even if he cannot pinpoint exactly who the individual is. An extension of k -anonymity, ℓ -diversity ensures that in any group of k individuals sharing the same quasi-identifiers, there are at least ℓ different values for the confidential attribute, preventing such attacks. t -closeness

⁸¹ *Ibid.*, p. 23-24.

⁸² *Ibid.*

⁸³ *Ibid.*

goes further, enforcing not only some diversity in the confidential attributes, but also that the confidential attributes in the groups of k individuals have the same distribution (up to a distance t) of the attribute in the whole dataset.

Probabilistic attacks are those in which access to a released dataset (either full access to an anonymised dataset or access through queries to the original dataset) changes the attacker's probabilistic belief (posterior probability in Bayesian statistics) on some attribute or attributes of her victim. In these attacks, the attacker does not directly learn the value of an attribute but can bound the probabilities that the attribute will take some specific value or values. ϵ -Differential privacy focuses on the prevention against such attacks.

4.2.3.2. Utility and privacy

Anonymisation mechanisms modify the original data to prevent disclosure of personal information, both of identities and their related confidential attributes. This modification may hurt the utility of the data. Anonymisation methods aim to protect the privacy of the respondents while producing as little as possible effects on the utility of the data. As utility, we may refer to the differences between the original and the modified data or to the differences in the results obtained from some processing on the original and the modified data (e.g., the accuracy of a machine learning model trained on the original data versus the accuracy of the model trained on the anonymised data).

Generic measures for data utility compare different statistics between the original dataset and the anonymised dataset. Some of these statistics include the means and covariances of some attribute subsets, and covariance and correlation matrices. Other measures include the mean absolute or mean squared errors between attributes in the original and anonymised datasets⁸⁴, or distance metrics (such as the Kullback-Leibler divergence) between the empirical distributions of the attributes⁸⁵. Propensity scores and cluster analysis indicate how likely it is to identify (or classify) the anonymised records as such when analysed together with the original records⁸⁵.

When the analyses to be performed on data are known before any anonymisation measure is applied, one can measure the utility of the anonymisation process by conducting such analyses on the original and the anonymised datasets and compare the results using some distance metric.

When trying to balance the utility of the anonymised datasets and the privacy guarantees offered by the anonymisation measures taken, one can follow two approaches:

⁸⁴ Domingo-Ferrer, J. and Torra, V., 2001. Disclosure control methods and information loss for microdata. Confidentiality, disclosure, and data access: theory and practical applications for statistical agencies, pp.91-110.

⁸⁵ Woo, M.J., Reiter, J.P., Oganian, A. and Karr, A.F., 2009. Global measures of data utility for microdata masked for disclosure limitation. Journal of Privacy and Confidentiality, 1(1).

- In utility-first anonymisation, the objective is to cause the minimum disruption to the dataset while minimising the disclosure risks. First, the data is transformed using some anonymisation method (as described in Section 4.2.3.4) with some parameters and then the identity or attribute disclosure risks are calculated. If the risk is considered too high, the anonymisation methods are re-run using more strict parameters. This process is repeated until the disclosure risks are considered low enough;
- In privacy-first anonymisation, some privacy model is enforced, such as k -anonymity or ϵ -differential privacy, that ensure some boundaries on the risks of re-identification or attribute disclosure. Anonymisation methods in this case are model dependent, and the parameters are derived from the model.

4.2.3.3. Privacy Models

k-Anonymity and extensions

A well-known privacy model is k -anonymity⁸⁶, which requires that each tuple of quasi-identifier attribute values be shared by at least k records in the database. This condition may be achieved through generalisation and suppression mechanisms, and through microaggregation⁸⁷.

Unfortunately, while this privacy model prevents identity disclosure, it may fail to protect against attribute disclosure. The definition of this privacy model establishes that complete re-identification is unfeasible within a group of records sharing the same tuple of perturbed quasi-identifier attribute values. However, if the records in the group have the same value (or very similar values) for a confidential attribute, the confidential attribute value of an individual linkable to the group is leaked.

To fix this problem, some extensions of k -anonymity have been proposed, the most popular being ℓ -diversity⁸⁸ and t -closeness⁸⁹. The property of ℓ -diversity is satisfied if there are at least ℓ 'well-represented' values for each confidential attribute in all groups sharing the values of the quasi-identifiers. The property of t -closeness is satisfied when the distance between the distribution of each confidential attribute within each group and the whole dataset is no more than a threshold t .

⁸⁶ Samarati, P. and Sweeney, L., 1998. Protecting privacy when disclosing information: k -anonymity and its enforcement through generalization and suppression.

⁸⁷ Domingo-Ferrer, J. and Torra, V., 2005. Ordinal, continuous, and heterogeneous k -anonymity through microaggregation. *Data Mining and Knowledge Discovery*, 11(2), pp.195-212.

⁸⁸ Machanavajjhala, A., Gehrke, J. and Kifer, D., 2006. ℓ -diversity: privacy beyond k -anonymity. *ICDE 2006. IEEE 22nd International Conference on Data Engineering*.

⁸⁹ Li, N., Li, T. and Venkatasubramanian, S., 2007, April. t -closeness: Privacy beyond k -anonymity and ℓ -diversity. In *2007 IEEE 23rd International Conference on Data Engineering* (pp. 106-115). IEEE.

ϵ -Differential Privacy and relaxations

Another important privacy model is differential privacy⁹⁰. This model was originally defined for queryable databases and consists in perturbing the original query result of a database before outputting it. This may be viewed as equivalent to perturbing the original data and then computing the queries over the modified data. Thus, differential privacy can also be seen as a privacy model for microdata sets.

An ϵ -differentially private algorithm is one that, when run on two datasets that differ in a single record, performs similarly (up to a power of ϵ) in both cases. That is, the presence or the absence of any single record does not significantly alter the output of the algorithm. Typically, ϵ -differential privacy is attained by adding Laplace noise with zero mean and parameter $\Delta(f)/\epsilon$, where $\Delta(f)$ is the sensitivity of the algorithm (the maximum change in the algorithm output that can be caused by a change in a single record in the absence of noise) and ϵ is a privacy parameter; the larger ϵ , the less privacy.

(ϵ, δ) -differential privacy allows for a small probability δ of data leaks. This parameter δ should be smaller than the inverse of the size of the dataset.

Rényi differential privacy⁹¹ is an alternative relaxation which does not allow for complete privacy compromises as (ϵ, δ) -differential privacy, leads to tighter privacy bounds and makes it easier to track the privacy budget, a property highly appreciated in federated learning scenarios.

4.2.3.4. Methods

In SDC, masking refers to the process of obtaining an anonymised dataset X' by modifying the original X . Masking can be perturbative or non-perturbative. In the former approach, the data values of X are perturbed to obtain X' . In contrast, in nonperturbative masking X' is obtained by removing some values and/or by making them more general; yet the information in X' is still true, although less detailed; as an example, a value might be replaced by a range containing the original value.

Perturbative masking

Perturbative masking generates a modified version of the microdata set such that the privacy of the respondents is protected to a certain extent while simultaneously some statistical properties of the data are preserved. Well-known perturbative masking methods include:

- Noise addition. This is the most popular method, which consists in adding a noise vector to each record in the dataset. The utility preservation depends on the amount and the distribution of the noise;

⁹⁰ Dwork, C. and Roth, A., 2014. The algorithmic foundations of differential privacy. *Found. Trends Theor. Comput. Sci.*, 9(3-4), pp.211-407.

⁹¹ Mironov, I., 2017, August. Rényi differential privacy. In 2017 IEEE 30th Computer Security Foundations Symposium (CSF) (pp. 263-275). IEEE.

- Data swapping. This technique exchanges the values of the attributes randomly among individual records. Clearly, univariate distributions are preserved, but multivariate distributions may be substantially harmed unless swaps of very different values are ruled out;
- Microaggregation. This groups similar records together and releases the average record of each group. The more similar the records in a group, the more data utility is preserved.

Non-perturbative masking

Common non-perturbative methods include:

- Sampling. Instead of publishing the whole dataset, only a sample of it is released;
- Generalisation. The values of the different attributes are recoded in new, more general categories such that the information remains the same, albeit less specific;
- Top/bottom coding. In line with the previous method, values above (resp. below) a certain threshold are grouped together into a single category;
- Local suppression. If a combination of quasi-identifier values is shared by too few records, it may lead to re-identification. This method relies on replacing certain individual attribute values with missing values, so that the number of records sharing a particular combination of quasi-identifier values becomes larger.

Synthetic microdata generation

An anonymisation approach alternative to masking is synthetic data generation. That is, instead of modifying the original data set, a simulated dataset is generated such that it preserves some properties of the original data set. The main advantage of synthetic data is that no respondent re-identification seems possible since the data are artificial. However, if, by chance, a synthetic record is very close to an original one, the respondent of the latter record will not feel safe when the former record is released. In addition, the utility of synthetic data sets is limited to preserving the statistical properties selected at the time of data synthesis.

Some examples of synthetic generation include methods based on multiple imputation⁹² and methods that preserve means and co-variances⁹³. An effective alternative to the drawbacks of purely synthetic data are hybrid data, which mix original and synthetic data and are therefore richer⁹⁴. Yet another alternative is partially synthetic data, whereby only the most sensitive original data values are replaced by synthetic values.

⁹² Rubin, D.B., 1993. Discussion statistical disclosure limitation. *Journal of official Statistics*, 9(2), p.461.

⁹³ Burrige, J., 2003. Information preserving statistical obfuscation. *Statistics and Computing*, 13(4), pp.321-327.

⁹⁴ Domingo-Ferrer, J. and González-Nicolás, Ú., 2010. Hybrid microdata using microaggregation. *Information Sciences*, 180(15), pp.2834-2844.

5. Privacy of Location-Based Services

The advancement of location sensing technologies as GPS or RFID at user level has made location-based services (LBS) ubiquitous. With smartphones as the primary medium, everyday millions of people use applications such as route planners (*e.g.*, to find the best route between two locations), social networking (*e.g.*, to find potential friends based on their location and profile) and local business search (*e.g.*, to find the nearest restaurant). Consequently, a huge amount of personal data is collected by service providers. This section discusses current privacy issues and solutions regarding this type of data, based on surveys by Chow *et al.*⁹⁵, Christin⁹⁶, and Decker⁹⁷ and a critical search.

This section is organised as follows. First, the unique characteristics of privacy in LBS are explained. Following that, the most common risks to this data are listed. A description of the most notorious privacy methods is made in the third subsection.

5.1. Unique characteristics

Determined by their use cases, location-based services have some peculiarities that are important to know in order to understand their privacy issues and the methods to provide protection against privacy invasions. In the following, the most notable features are defined.

The main dissimilarity of LBS with the rest of mobility data frameworks is that most LBS operate as real-time systems, having to process a growing number of locations and queries in a limited time. Consequently, the methods proposed for other frameworks are often incompatible, as they are designed for historical data.

Another feature that should not be ignored is the requirement for a minimum level of consistency in identification. In this case, identification refers to the value that allows the LBS provider to know who the user is, not necessarily a personal identifier of the user such as the name. Typically, this is a pseudonym, or the query code used to know who made the communication. On this basis, an LBS requires user identification to have one of the following types of consistency:

⁹⁵ Chow, C.Y. and Mokbel, M.F., 2011. Trajectory privacy in location-based services and data publication. ACM Sigkdd Explorations Newsletter, 13(1), pp.19-29.

⁹⁶ Christin, D., 2016. Privacy in mobile participatory sensing: Current trends and future challenges. Journal of Systems and Software, 116, pp.57-68.

⁹⁷ Decker, M., 2008, July. Location privacy-an overview. In 2008 7th International Conference on Mobile Business (pp. 221-230). IEEE.

- **Consistent:** The user's identifier does not change during their trajectory. Also known as trajectory data, it is common when the service provided requires user-specific information (e.g., "Recommend me nearby restaurants based on my profile") or the communication has a consistent and periodic query (e.g., "Continuously tell me the nearest shopping centre"). In the second case it is the query with its parameters which is considered the identifier since it does not change during the trajectory;
- **Partially consistent:** The user's identifier changes at some points along the path. Consequently, in some parts of the trajectory the identification remains consistent, until the identifier changes. This type of data results from some privacy methods and can be used for mobile participatory sensing;
- **Inconsistent:** The user's identifier changes at each position of the trajectory. Also known as positional data, it is common when services respond to location-based queries without requiring user-specific information (e.g., "Show me the offers of the nearest shops").

Note that this is a hierarchical classification where, for example, any system requiring inconsistent identification can operate using data with consistent identifiers, but not vice versa.

Finally, LBS are particularly sensitive to data suppression or ignoring. This is because, in the context of an application, suppressing or ignoring many location-based queries of a specific user would be equivalent to denying their service, significantly affecting the user experience.

5.2. Risks

It is well-known that mobility data can cause serious privacy problems if not properly processed, and LBS data is not an exception. Moreover, LBS have particularities that affect privacy in ways that need to be mentioned. In this section these LBS-specific risks will be explained, together with a definition of common privacy issues in terms of identification consistency.

The most distinctive risk of LBS is the publication of fully identified sensitive data (e.g., posting your home's location in a friends' finder application). Although it is assumed that the user consciously publishes this sensitive information and agrees to it, they often do not realise what the consequences could be. In addition, users are sometimes unaware of when and how they are sharing this data. As an illustration, periodically posting their afternoon running route may seem safe to them, but if it always starts and/or ends at the door of their house, home identification would be very easy for an attacker.

On the other hand, regardless of the information sensitivity, fully identified location data could be an extremely useful resource for an attacker, as it allows for record linking to a secondary anonymised dataset which does contain sensitive data. For example, if the user posts their location in a popular public event (non-sensitive data) but it is the only location linked to that user in such a secondary dataset, re-identification is possible.

Additionally, it is important to highlight that the LBS data is often related to personal preferences (*i.e.*, queries to applications that recommend food, shopping centres or cafés). As these preferences are often shared on social networks (e.g., with photos of activities performed), record linkage is more likely to occur than for other types of mobility data.

Finally, for a broader understanding, common privacy risks for each type of identification consistency are defined below:

- **Consistent:** Since the path information is complete, it is the most valuable type of data for an attacker. If it is not properly anonymised, re-identification can be performed with a small number of random user positions. It can be used to obtain highly sensitive information as home or work location and mobility patterns;
- **Partially consistent:** Re-identification is harder than for consistent identification and, as the attacker only re-identifies a part of the trajectory, the obtained information is probably less valuable. However, there are time- and velocity-based attacks that may allow the attacker to transform the data to one with consistent identification;
- **Inconsistent:** This is the least valuable type data for an attacker, as it is difficult to acquire new user-specific information from it. Nevertheless, there are target tracking techniques that can obtain the user's trajectory, resulting in consistently identified data.

5.3. Privacy methods

In order to reduce the privacy risks of the LBS, several methods have been proposed in the literature in recent years, following the growing popularity of these services. Like methods for other types of mobility data, these have anonymisation as focus. Nevertheless, because of the unique characteristics of LBS, the proposed approaches present dissimilarities with those used for other mobility data. Consequently, this subsection is structured as follows. First, an explanation of differentiating features of these methods is provided. Next, the most notorious privacy methods for LBS are presented. Finally, a summary table of the presented methods is shown.

5.3.1. Differentiating features

A first important feature of LBS is that most of these services operate in real time. Consequently, the corresponding methods must be able to anonymise growing data with limited time costs. This is contrary to many mobility data anonymisation techniques, where historical data is expected, and time is not critical. On this basis, all the presented methods are designed to work in real time.

Another characteristic to consider is that the LBS server is considered as untrustworthy. Therefore, anonymisation must be done beforehand, requiring a Trusted Third Party (TTP) or a Peer-to-Peer (P2P) network between users to perform the process and then forward the data to the service provider. The use of TTPs has become commonplace in recent years, however, as a TTP is a single point of failure, having a truly trustworthy server can be complex and costly. Therefore, from a privacy and economic point of view, a P2P-based approach is preferable.

Finally, LBS privacy methods differ from others in their sensitivity to data suppression. The reason for this difference is that many of the other methods enhance privacy by suppressing location data that is difficult to anonymise, but in an application, this could lead to denial of service for some users. The proposed methods try to maximise the number of anonymised locations.

5.3.2. Methods

In the following, the most notorious privacy methods are briefly defined and discussed. The order is based on the identification consistency they work with and whether they are complete privacy methods on their own or just additional techniques that can be used in conjunction with any of the other methods.

It is important to point out that many of the methods presented are based on k -anonymity and, consequently, may suffer from a homogeneity attack. This is especially likely when spatial cloaking is used without a minimum area requirement. This drawback will not be mentioned for each concerned method in order not to be repetitive, but it should be considered.

5.3.2.1. Group-based

Proposed by C. Y. Chow and M. F. Mokbel⁹⁸, this method aims to protect real-time location trajectory queries with consistent identifiers via spatial cloaking.

The idea is to obtain k -anonymity, grouping nearby users at the start of their trajectory, and always reporting the region containing the k users as the location. To explain it in more detail, the steps of the method are listed below, followed by Figure 1 depicting an example of 3-anonymity with six users.

1. The anonymisation system (placed on a TTP server) receives the first position of the users' trajectories.
2. The method creates groups of k nearby users. This groups are intended to be not modified during the trajectories.
3. For each group, a rectangular minimum bounding region (MBR) is created (Figure 1a) based on their locations. This rectangle is the cloaking region of the group, as it achieves the corresponding k -anonymisation. The region is sent to the LBS provider as the location of each user of the group.
4. Users continue sending their position to the anonymisation server until the end of their path, repeating the third step.

⁹⁸ Chow, C.Y. and Mokbel, M.F., 2007, July. Enabling private continuous queries for revealed user locations. In International Symposium on Spatial and Temporal Databases (pp. 258-275). Springer, Berlin, Heidelberg.

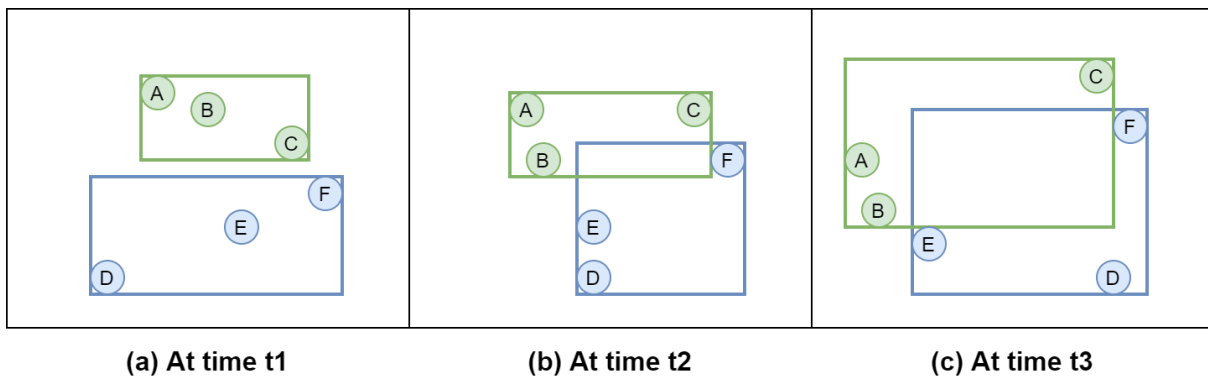


Figure 1: Example of the Group-based approach with six users. The two groups are coloured green and blue respectively. Each user is represented by a circle with the pseudonym inside. Rectangles are used to depict the cloaking regions.

This proposal has some privacy issues that should be mentioned. The first problem comes from the edges and corners of the region. If the user's location has x and y coordinates that are the maximum or minimum of the group, the position will be equal to one of the corners of the region. An example of this can be seen in Figure 1, where the top-right corner of the blue region is always equal to the position of the user F . This problem is also present at times t_2 and t_3 for the user C and the top-right corner of the green region. Another privacy risk of this approach comes from disconnections. If one or more users in a group disconnect, the k -anonymity will be broken, since the quantity of possible pseudonyms for a user in the group will be k minus the number of disconnected users.

On top of that, this method has a remarkable quality-of-service drawback. To put it briefly, the cloaked region tends to become very large after a long period of time. The next privacy method will address this issue.

Lastly, it's important to note that this method works under the assumption that multiple users start their trajectory at the same time and in close proximity and end at the same time, which is unlikely unless there are many users.

5.3.2.2. Distortion-based

This method, proposed by X. Pan, X. Meng, and J. Xu⁹⁹, is based on the group-based approach and aims to reduce the problem of very large cloaking regions.

To this end, the anonymiser requires the velocity of each user at the start of the trajectory (usually the velocity towards the target position) and assumes that it will be constant. Using this velocity, the

⁹⁹ Pan, X., Meng, X. and Xu, J., 2009, November. Distortion-based anonymity for continuous queries in location-based mobile services. In Proceedings of the 17th ACM SIGSPATIAL International Conference on Advances in Geographic Information Systems (pp. 256-265).

proposed method defines a distortion function that predicts the expected size of the cloaking region over time. When creating the groups, sets are sought that minimise the distortion function along the trajectory. In addition, the constant velocity assumption is used to predict the next cloaking regions.

Equivalent to the description of the previous method, the steps of this method are listed below along with an example (Figure 2). This example shows what the result of this approach would look like with the same data as the previous one.

1. The anonymisation system (located on a TTP server) receives the first position and velocity (including speed and direction) of a user.
2. Using the positions and velocities of the trajectories currently starting, the best groups that minimise the result of the distortion function are searched for.
3. The cloaking regions of the user's group are calculated and sent to the LBS provider.
4. The next cloaking region of the group is pre-computed assuming constant velocities.
5. All the members of the user's group submit the next location. Then, the predicted cloaking region is checked. If it does not contain all the locations, it is recomputed. Finally, the resulting cloaking region is sent to the LBS provider.
6. Users continue sending their position to the anonymisation server until the end of their path, repeating steps 4 and 5.

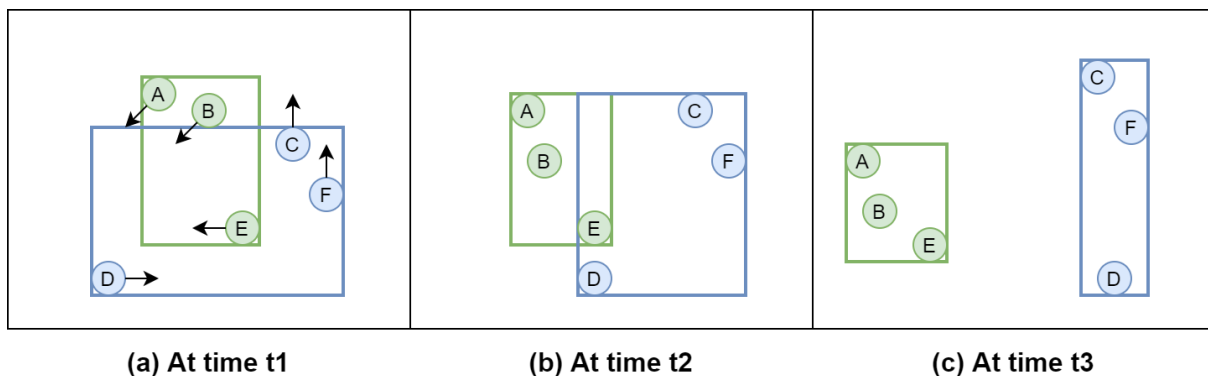


Figure 2: Example of the Distortion-based method with six users. The two groups are differentiated by colour. Each user is represented by a circle with a character as pseudonym and an arrow depicting their expected movement direction. Rectangular regions are used to illustrate the cloaking areas.

While this method improves the utility of the data with respect to the previous one, it's important to note that it still carries over the privacy issues from the previous proposal.

5.3.2.3. Mix-zones

Proposed by A. R. Beresford, and F. Stajano¹⁰⁰, mix-zones performs real-time anonymisation using the mix-networks as an inspiration, transforming data with consistent identification into data with partially consistent identification.

Mix-networks are networks with normal routers and mix-routers. A mix-router collects k packets of the same size inputs and outputs them in a random order, thus ensuring unlinkability between incoming and outgoing messages. This concept has been translated into LBS with mix-zones as equivalent to the mix-routers and user data as analogous to messages. Mix-zones are limited areas that are usually located around sensitive places such as hospitals or religious buildings. The user data consist of the users' locations and current pseudonyms. The steps followed by the method for each received location are explained below, together with Figure 3 for exemplification.

1. The anonymisation server (TTP) receives the location of a user.
2. If the user is inside a mix-zone, the location is not forwarded to the LBS.
3. Otherwise, if the user is not in a mix-zone, the system checks if the user has just exited a mix-zone. If true, their pseudonym is changed to a new, unused one. The location and corresponding pseudonym are sent to the LBS.




Figure 3: Example of the Mix-zones method with three users. Each pseudonym's trajectory has a different colour, a start and end position defined by a circle with the pseudonym inside and an arrow defining the trajectory. The Mix-zone region is defined by a white rectangle. The user's pseudonyms have the following changes: $A \rightarrow F$, $B \rightarrow D$, $C \rightarrow E$.

This approach will give k -anonymity to a set of users under the following requirements:

¹⁰⁰ Beresford, A.R. and Stajano, F., 2003. Location privacy in pervasive computing. IEEE Pervasive computing, 2(1), pp.46-55.

- The set has at least k users;
- The users are in the same mix-zone;
- Each user spends a random amount of time inside the mix-zone;
- The probability of a user entering through an entry point is the same as the probability of exiting through any of the exit points.

Note that in some cases, it is difficult to meet all these requirements. For example, if users move with a constant speed from one side of the mix-zone to the other, their time inside will not be random, which enables a time-based attack using the resulting first in first out behaviour. As another example, if the mix-zone is located on an avenue, the entry and exit points of the mix-zone will not be random, allowing the attacker to know that a person entering at one side is likely to exit from the other. These problems are extremely common for vehicular mix-zones, as defined in Julien *et al.*¹⁰¹ and Palanisamy *et al.*¹⁰², where a specific distribution of mix-zones is defined to try to solve them.

It is also important to mention that this method does not have a well-defined level of privacy. That is, there is not a parameter that allows for a comprehensible definition of the trade-off between privacy and utility. This sets it apart from other privacy methods, where one or more parameters can be set for that purpose. The most well-known example is the k parameter of k -anonymity, which clearly defines the privacy level. This issue can lead to usability problems, as a certain privacy cannot be guaranteed to the user.

Finally, a major problem with this method is the mix-zones placement, as the quantity and location will directly affect privacy. Research has been done to optimise this placement problem¹⁰³ but, as will be seen in the next method, this is not the only approach.

¹⁰¹ Julien, F., Raya, M., Felegyhazi, M. and Papadimitratos, P., 2007. Mixzones for location privacy in vehicular networks. In Association for Computing Machinery (ACM) Workshop on Wireless Networking for Intelligent Transportation Systems (WiN-ITS).

¹⁰² Palanisamy, B. and Liu, L., 2011, April. Mobimix: Protecting location privacy with mix-zones over road networks. In 2011 IEEE 27th International Conference on Data Engineering (pp. 494-505). IEEE.

¹⁰³ Freudiger, J., Shokri, R. and Hubaux, J.P., 2009, August. On the optimal placement of mix zones. In International Symposium on Privacy Enhancing Technologies Symposium (pp. 216-234). Springer, Berlin, Heidelberg.

5.3.2.4. SwapMob

SwapMob, proposed by J. Salas, D. Megías, and V. Torra¹⁰⁴, is a method based on mix-zones whose main contribution is the non-restriction of the mixing space, eliminating the mix-zones placement problem.

The concept, briefly summarised, is to mimic the behaviour of a mix-zone whenever two or more users' trajectories cross with each other, rather than in fixed mix-zones. In the following, the method process will be explained in more detail.

Users periodically send their location to the anonymisation server, located on a TTP. If the anonymiser detects that two or more users are close to each other (using predefined temporal and spatial thresholds) their pseudonyms are randomly swapped. Pseudonyms are exchanged instead of being replaced by new ones so that the attacker cannot easily know when the mix has occurred. Finally, the anonymiser sends the locations with the corresponding pseudonyms to the LBS provider. Note that, in this context, the exchange of pseudonyms is equivalent to the exchange of partial paths. For ease of comprehension, Figure 4 depicts an example of this method.




Figure 4: Example of SwapMob with three users. Each pseudonym's trajectory has a different colour, an initial and final position defined by a circle with the pseudonym inside and an arrow which defining the path.

Although this approach solves the mix-zones placement problem, it retains some drawbacks. The main one is that it still lacks a well-defined level of privacy. Privacy mainly depends on the users' trajectories, as it defines the quantity and location of swaps. Additionally, there is still the possibility of a velocity-based attack. As an example, if a swap is performed between two trajectories which intersect at a large angle (e.g., the swap at the 90-degrees intersection of A and B in Figure 4b), it

¹⁰⁴ Salas, J., Megías, D. and Torra, V., 2018, September. SwapMob: Swapping trajectories for mobility anonymization. In International Conference on Privacy in Statistical Databases (pp. 331-346). Springer, Cham.

is easy for the attacker to notice the swap (the direction has dramatically changed) and know exactly the exchanged pseudonyms.

5.3.2.5. P2P spatial cloaking

Proposed by C. Chow, M. F. Mokbel, and X. Liu¹⁰⁵, this method performs real-time personalised anonymisation for data with inconsistent identifiers using a P2P network.

The core idea is to use the P2P network to search nearby users and to realise spatial cloaking without requiring a TTP. Additionally, each user has a so-called privacy profile in which they can define the desired k for k -anonymity and the minimum area for cloaking (as an approximation to ℓ -diversity). In the following, the simpler version of this approach will be defined together with the example in Figure 5. Subsequently, some extensions that improve the privacy and utility of this approach are listed.

When a user requests to send a location to the LBS provider, the following anonymisation process is initiated:

1. The user device uses the P2P network to search $k - 1$ nearby users. In order to do so, the user broadcast a request to the peers that are one hop from itself.
2. If the number of neighbour users at the current distance (initially one hop) is less than $k - 1$, the search is extended. In other words, the broadcast is repeated increasing the number of hops (performing a multi-hop request) until at least $k - 1$ users' locations are found. If the required number of locations cannot be obtained, this step can be repeated periodically or k can be temporally reduced.
3. Once $k - 1$ or more users' locations have been found, the algorithm selects the $k - 1$ nearest locations and checks if the resulting cloaking region has the minimum area required by the user. If it not, the region is expanded in width and length to reach the desired minimum area.
4. The resulting cloaking region is sent to the LBS provider as representation of the current location.

¹⁰⁵ Chow, C.Y., Mokbel, M.F. and Liu, X., 2006, November. A peer-to-peer spatial cloaking algorithm for anonymous location-based service. In Proceedings of the 14th annual ACM international symposium on Advances in geographic information systems (pp. 171-178).




Figure 5: Example of P2P spatial cloaking for the user A. On the left side, the P2P network is depicted, using edges to represent connections between nodes. Blue nodes are one hop away from A, and those two hops away are coloured purple. The right-hand side of the figure represents the resulting cloaking region of A for a k of four.

The above-mentioned method has multiple potential problems. Consequently, the authors presented the following extensions:

- *Information sharing scheme and historical sharing scheme:* Due to scarce communication resources in mobile P2P environments, excessively scanning the network for peers could pose a scalability issue. On the other hand, since we consider a highly *ad hoc* mobile environment, where neither fixed communication infrastructure nor centralised/distributed servers are available, mobile users can only communicate with each other via multi-hop peer-to-peer routing. Due to user mobility, users can be partitioned into disjoint networks. When a network partition takes place, a mobile user can only communicate with other peers residing in their network partition. Therefore, if a user's required anonymity level is greater than the number of users residing in the network partition, the user cannot find enough peer location information to blur their position; and thus, the user suffers from a network partition problem. To reduce both problems, a cache-based system is proposed. Each user's device stores the data obtained from the last request with the corresponding timestamp. Then, when location data is required by the user's device or a neighbour, this cache can be used instead of the broadcasting process. Thus, the information sharing scheme can reduce communication and computational overhead. Note that the cache can be used if and only if the time elapsed is not more than a predefined threshold, which ensures that the locations are not too old;
- *Cloaked area adjustment scheme:* As defined above, the $k - 1$ closest locations are selected to compute the cloaking area of a user. This peer selection process can lead to a privacy breach, as the user's location tends to be the one closest to the centre of the region. In other words, a centre-of-cloaked-area privacy attack can be very successful. To prevent this problem, the region must be modified so that the probability that the user's location is the closest to the centre is $1/k$. For that, a cloaked area adjustment algorithm is computed.

As has been seen, this is a simple approach that provides k -anonymity using a P2P system. Additionally, some extensions have been presented to increase the utility and privacy. However, there are some privacy issues to highlight. The first one is that users send their location to the LBS provider directly from their device. Therefore, if their IPs do not vary, these could be used to track them, changing identification consistency from inconsistent to consistent. On the other hand, this

method has no user control, allowing malicious users to query the locations of nearby users in the P2P network as much as they want.

5.3.2.6. P⁴QS

The Peer-to-Peer Privacy Preserving Query Service (P⁴QS)¹⁰⁶ is an approach proposed by M. Ghaffari, N. Ghadiri, M. H. Manshaei, and M. S. Lahijani for the anonymisation of real time data with inconsistent identification using P2P networks.

This method aims to provide a P2P-based architecture which does not allow either LBS provider or malicious users to obtain private data. Therefore, it addresses the privacy concerns of the aforementioned P2P spatial cloaking method, at the cost of complexity. In the following, the approach will be explained in two parts. In the first one, the architecture will be presented, including the different roles of the peers. The second part will define the protocol based on these definitions.

P⁴QS has the LBS provider and the P2P peers (users) as main actors. The LBS provider is located on a public server and owns some signing key. In addition, it has the IPs of the users and can send them *tickets*. A *ticket* is a packet of data that users' queries must contain in order to be attended by the LBS. Each ticket is signed with server's signing key and contains a random token, an expiry time, and an identifier for validation. Thus, each entity in the system can verify the tickets' validity with the server's public key and check the identifier (and expiry time) to validate it. On the other hand, peers have a random identifier called MyRA and are part of a Distributed Hash Table (DHT) along with the other peers. The DHT allows peers to find other peers based on their MyRA. Each peer can play one or more of the following roles at the same time:

- Client: Every peer of the network is potentially a client. It sends location-based queries to an anonymising peer and waits for the response from the LBS to the corresponding broker peer;
- Anonymiser: It is responsible for the anonymisation of queries in an area defined by its MyRA. It receives the queries from the clients (along with other information), anonymises the location and sends the resulting message to the LBS provider. As this role depends on the MyRA identifier, it is fixed until disconnection;
- Broker: It receives the response to a client's query and waits for the client to request for their response. Then, it sends the packet to the client. This role is defined by the peer's MyRA and the ticket used by the client.

Based on the above definitions, the initialisation and the main loop of the protocol are defined below.

¹⁰⁶ Ghaffari, M., Ghadiri, N., Manshaei, M.H. and Lahijani, M.S., 2017. P⁴QS: A peer-to-peer privacy preserving query service for location-based mobile applications. *IEEE Transactions on Vehicular Technology*, 66(10), pp.9458-9469.

First, when the system starts, it performs the following steps:

1. Each peer initialises its own MyRA and joins the P2P network. The DHT of the peers is created with MyRA as key/identifier of each peer.
2. The LBS sends the first set of tickets to the peers. The peers will exchange their tickets periodically over time, so that the LBS cannot know who the user is based on their tickets. The LBS server will send new tickets in predefined time periods.

When a client peer wants to submit a query, the following steps are performed:

1. A valid ticket is selected from its ticket list.
2. The anonymiser of the zone is found using the DHT. This will be the peer with the MyRA closest to the hash of the current location. A custom hash which produces similar values to close locations is used.
3. The broker of the response is found using the DHT. It will be the peer with the MyRA closest to the hash of the ticket. Then, a request for response is sent, letting the broker know that the future response is for this client.
4. The client sends the location, query, ticket, broker's IP, and a proposed key to the anonymiser. The query, ticket and IP are encrypted with the proposed key (symmetric cryptosystem). The proposed key is encrypted with the server's public key. Thus, only the LBS server can access the encrypted fields.
5. When the anonymiser receives the message, it tries to k -anonymise the location using spatial and temporal cloaking. If, after a timeout, it fails to do so, it tries to collaborate with adjacent anonymisers or generate fake queries (using valid tickets). Once k queries are anonymised (ignoring if some of them are fake) they are sent to the LBS provider.
6. When the LBS receives a query, it decrypts the proposed key using the server's private key and uses this one to decrypt the rest of the data. Then, if the ticket is valid, the query is processed, the response is encrypted with the client's proposed key and the response packet is sent to the broker's IP. Otherwise, the query is ignored.
7. When the broker receives the response, it checks whether a request for it has been received. If there is a request, the response is sent to the corresponding client. Otherwise, the response is stored until the request is received.
8. The client decrypts the received response using its proposed key.

Now that the method has been defined, the privacy and utility issues related to each actor in the architecture are listed below:

- Client: Their identity is kept anonymous from the server; therefore, their trajectory is also protected;
- Anonymiser: It has access to the identity and exact positions of the clients. Nevertheless, as its role is randomly selected and it only obtains data from clients in a specific area, the probability

that a malicious anonymiser that can perform client identification and obtain new knowledge is low.

On the other hand, anonymisers are points of failure that could be overloaded, denying service in the corresponding area. For example, in case the anonymiser is assigned to a crowded zone. To avoid this problem, the authors propose that the peer disconnects and reconnects to the system, obtaining a new MyRA. Thus, its load will be divided between the two closest peers in the DHT;

- LBS provider: It receives the queries from the anonymisers and sends the response to the brokers, so the only way to identify the client is through the ticket. However, as the tickets are exchanged randomly between peers, this correlation is unlikely. Moreover, if a malicious peer wants to make a DoS attack on the system, they will need to send plenty of queries with valid tickets. Therefore, as valid tickets are time-limited the DoS attack can be controlled. However, the server cannot identify the malicious peer;
- Broker: It receives encrypted responses from LBS server depending on the hash of the ticket used and the DHT. Therefore, it can't obtain new knowledge.

Finally, it is important to highlight a privacy problem that, despite having a potentially simple solution, can be dangerous. The problem is that no minimum area for cloaking is defined, allowing regions to be as small as the distance between users, which can be small enough to compromise user's privacy.

5.3.2.7. CliqueCloak

Suggested by B. Gedik, and L. Liu¹⁰⁷, CliqueCloak is a method that uses spatial cloaking and graph theory to anonymise data with inconsistent identifiers in real time, allowing for user-personalised privacy.

The goal of this method, like others, is to obtain k -anonymity using the proximity between users and based on users' privacy profiles. These profiles define the minimum k , maximum temporal tolerance and maximum spatial tolerance allowed by each user. The system will attempt to anonymise the users' locations meeting all these requirements. For this purpose, a constraint graph is used. For ease of understanding, the steps followed by the method are listed together with an example (Figure 6).

1. The anonymiser (running on a TTP) receives messages from the users, including their location, privacy profile and message identifier.

¹⁰⁷ Gedik, B. and Liu, L., 2004. A customizable k-anonymity model for protecting location privacy. Georgia Institute of Technology.

2. Using the temporal tolerance desired by the users, deadlines are set. If the deadline of a message is reached without being anonymised, it is said that the privacy requirements cannot be accomplished and then the message is dropped.
3. With the spatial tolerance of the privacy profile, a square region is established around the location of each user. If other users are inside this region, they are related. This is represented as an edge between the corresponding users' nodes of the constraint graph.
4. Based on the previous information, the method searches for a set of related positions in the constraint graph that archives the required k -anonymity for the maximum number of users.
5. Once the group is created, the minimum bounding rectangle (MBR), also called cloak region, is computed as the minimum rectangle that contains all directly related positions of the group.
6. Each user's message is sent to the LBS provider with the cloak region of their group as the position.

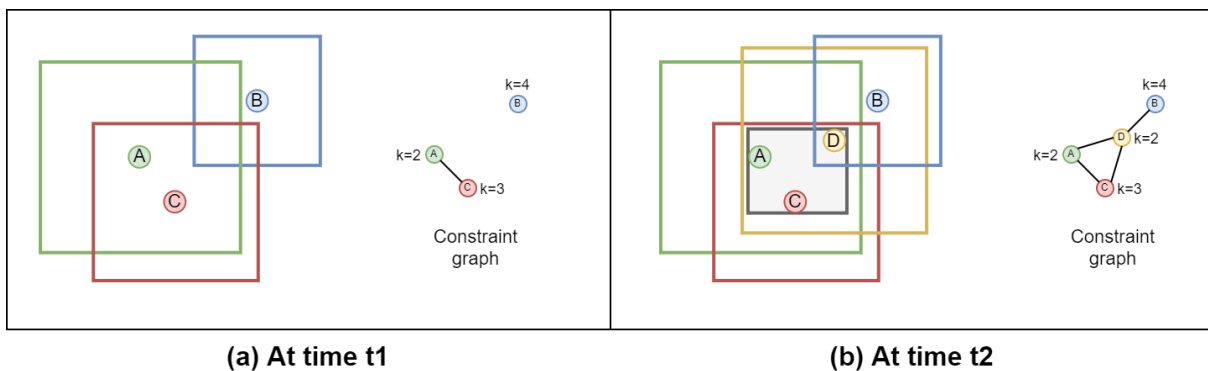


Figure 6: *Example of CliqueCloak with three users at time t_1 and four users at time t_2 . Users are represented as coloured circles with the character of the user's pseudonym. The maximum spatial tolerance is defined by squares centred in the corresponding user. The current constraint graph is shown on the right. The user's desired k is displayed next to the respective node of the constraint graph. The resulting cloaking region is represented as a grey rectangle.*

Note that in Figure 6a, it is not possible to create a group that meets all user requirements. In that case, only A would be anonymised by creating a group with C , ignoring B and C . However, in Figure 6b, the location of user D arrives, allowing all locations to be anonymised.

Finally, it is important to highlight a privacy problem that, despite having a potentially simple solution, can be dangerous. The problem is that no minimum area for cloaking is defined, allowing regions to be as small as the distance between users, which can be small enough to compromise user's privacy.

5.3.2.8. PrivacyGrid

PrivacyGrid is a method proposed by B. Bamba and L. Liu¹⁰⁸ which uses regular spatial decomposition to protect real-time data with inconsistent identifiers, integrating a user-personalised privacy.

This method has multiple similarities with CliqueCloak, as it also anonymises data in real time with inconsistent identifiers, allows each user to have their own privacy profile, uses a TTP and searches for the minimum cloaking region that meets the privacy requirements of the users. The main difference is in how the algorithm finds the cloaking region. The core of the idea is to limit the working area (e.g., a city) to a rectangle and divide it into cells of regular size, forming a grid. For each cell of this grid, the system will have an object count, that is, how many users are currently inside. Additionally, the current cell of each user is stored. When the anonymiser server receives a message from a user, the next steps are followed:

1. The location of the user is mapped to the grid space, obtaining the current cell;
2. If the current cell is different from the previous one, the users' current cell and the grid are updated accordingly;
3. The current cell is selected as a first attempt to find the cloaking region;
4. If the object count of the region is equal to or greater than the desired k of the user's privacy profile, the masking region has been found. Otherwise, this step is repeated by adding the adjacent cells. If the privacy requirement is not met despite using all cells in the grid, the message is said to be non-anonymisable and the process is terminated;
5. If the cloaking region is found, it is sent to the LBS provider together with a pseudonym (different for each message).

Note that the use of the grid is intended to reduce processing time as with it, it is not necessary to check the location of each current user to find near users, but only the current and adjacent cells in the grid.

¹⁰⁸ Bamba, B., Liu, L., Pesti, P. and Wang, T., 2008, April. Supporting anonymous location queries in mobile environments with privacygrid. In Proceedings of the 17th international conference on World Wide Web (pp. 237-246).

5.3.2.9. Path Confusion

This method proposed by B. Hoh, M. Gruteser, H. Xiong, and A. Alrabady¹⁰⁹ aims to protect real-time position and velocity data with inconsistent identifiers against trajectory disclosure using an uncertainty-aware path cloaking algorithm. It tries to offer a higher accuracy than other similar methods such as CliqueCloak or PrivacyGrid.

The core concepts of this method are the target tracking technique and the Time-To-Confusion metric. Target tracking is defined as a methodology which finds the path followed by a user in the inconsistently identified data. For ease of comprehension, the steps followed by the simplest version of target tracking are listed below:

1. Given a user's sample at time 1 with the corresponding position and velocity, the position at time 2 is predicted. For that, the fixed elapsed time between samples is used and the velocity is assumed to be constant.
2. The distance between the predicted position at time 2 is calculated, since the closest one is the most likely to be the user's next position. However, knowing that predicted position may be incorrect, a further filter is performed. The k closest samples are selected, and an information-theoretic uncertainty metric is computed. If the uncertainty is less than a parameter called ConfusionLevel, the closest sample is considered the user's sample at time 2. Otherwise, the track is lost.
3. If the track is not lost, the process is repeated using the sample predicted for time 2 as the sample for time 1.

On the other hand, the Time-To-Confusion metric is described as the time a user can be tracked. This is computed as the time elapsed between the initial sample and the last sample tracked by the target tracking technique. The approach uses the Time-To-Confusion metric as the privacy metric, considering it as inverse to the user's privacy. Concretely, a maximum Time-To-Confusion needs to be defined, together with the parameters of the target tracking procedure (k number of samples to check and maximum confusion level accepted).

From the above definitions, the privacy-preserving mechanism of the method's trajectory can be defined. First, the anonymisation system (running on a TTP server) receives a position and velocity data of a user. Then, the system applies the tracking technique to check whether this sample allows to correctly track the user for more than the specified maximum. If the Time-To-Confusion is not exceeded, this sample will be released to the service provider. Otherwise, it will be ignored.

¹⁰⁹ Hoh, B., Gruteser, M., Xiong, H. and Alrabady, A., 2010. Achieving guaranteed anonymity in gps traces via uncertainty-aware path cloaking. IEEE Transactions on Mobile Computing, 9(8), pp.1089-1107.

While this approach has been proven to provide better data accuracy than CliqueCloak (and probably PrivacyGrid), it is not perfect. In particular, the assumption of Time-To-Confusion as privacy metric has an issue that needs to be mentioned. The problem is that even if the system perfectly provides a reduced Time-To-Confusion for a user (e.g., five minutes), the non-confused section of the path may provide sensitive information to an attacker. To address this problem, the authors proposed some extensions to the algorithm. These extensions aim to maximise the protection for sensitive origins (by not initiating the releasing until the first confusion point is reached) and destinations (by waiting before releasing the data to avoid publishing the final locations). Nevertheless, it is difficult to know whether these extensions prevent the attacker from obtaining any sensitive data.

5.3.2.10. Silent Zones

Suggested by K. Wiesner, S. Feld, F. Dorfmeister, and C. Linnhoff-Popien¹¹⁰, Silent Zones is an additional method (can be used together with any of the other methods) which defines regions around sensitive locations where mobility data will not be shared. It is an evolution of the so-called Ban Zones¹¹¹ proposed by J. Krumm, with the main difference that the regions have different sizes depending on the buildings surrounding the sensitive place.

Specifically, this approach allows users to define sensitive locations (although these could also be predefined by the system) where they do not want to send their mobility data to the service provider. Then, when the user is inside of one of these Silent Zones the method computes the minimum region containing it and k buildings, guaranteeing k -anonymity in terms of surrounding buildings. This technique can be performed in their own device or a TTP server, and it would require from local mapping information to know where the buildings are.

5.3.2.11. Reduction of precision

Reduction of precision is an additional approach (can be used together with any of the other methods) which applies spatial and temporal generalisation to increase privacy. Even though this cannot be used for services which require the highest possible accuracy, there are many others that do not, such as weather forecast, advertising, or location-based news. Also, it is possible that the

¹¹⁰ Wiesner, K., Feld, S., Dorfmeister, F. and Linnhoff-Popien, C., 2014, April. Right to silence: Establishing map-based silent zones for participatory sensing. In 2014 IEEE Ninth International Conference on Intelligent Sensors, Sensor Networks, and Information Processing (ISSNIP) (pp. 1-6). IEEE.

¹¹¹ Krumm, J., 2007, May. Inference attacks on location tracks. In International Conference on Pervasive Computing (pp. 127-143). Springer, Berlin, Heidelberg.

spatial and temporal precision requirements differ, allowing the resolution of only one of them to be reduced. Furthermore, it is important to note that this method can be performed on users' devices (without requiring any communication) or on a TTP server.

The simplest ways to perform precision reduction are truncation and rounding. An example of truncation for spatial data would be to eliminate the N last decimals of the coordinates. On the other side, if applied to temporal information, one of the units (e.g., seconds or minutes) can be removed. However, as M. Decker mentions, these techniques have a drawback when used for real-time LBS: the value jumps¹¹². For example, if spatial data truncation is performed and a user crosses the line between "2.9" and "3.0", an attacker will know it, as the resulting value will just change from "2" to "3". Consequently, in this case, the attacker would have information as if one less decimal had been removed. This drawback also affects rounding (e.g., when the value changes from "2.4" to "2.5") and temporal data.

An alternative to rounding and truncation is the application of random noise to the samples, which has a higher computational cost but avoids the drawbacks of the two previous techniques. An example of this method was proposed by J. Krumm¹¹³, where 2D Gaussian noise is added to the locations. Specifically, a sample-specific perturbation is performed by first choosing a random uniform direction and then a Gaussian-distributed magnitude using a predefined deviation (e.g., 50 metres).

5.3.3. Summary table

The following table summarises the main characteristics of the privacy methods presented.

Table 3: Main characteristics of the privacy methods at a glance

Method's name	Identification consistency	Where is performed?	Complete or additional?
Group-based	Consistent	TTP	Complete
Distortion-based	Consistent	TTP	Complete
Mix-zones	Partially consistent	TTP	Complete
SwapMob	Partially consistent	TTP	Complete
P2P spatial cloaking	Inconsistent	P2P	Complete
P ⁴ QS	Inconsistent	P2P	Complete

¹¹² Decker, M., 2008, July. Location privacy-an overview. In 2008 7th International Conference on Mobile Business (pp. 221-230). IEEE.

¹¹³ Krumm, J., 2007, May. Inference attacks on location tracks. In International Conference on Pervasive Computing (pp. 127-143). Springer, Berlin, Heidelberg.

CliqueCloak	Inconsistent	TTP	Complete
PrivacyGrid	Inconsistent	TTP	Complete
Path Confusion	Inconsistent	TTP	Complete
Silent Zones	Any	User's device or TTP	Additional
Reduction of precision	Any	User's device or TTP	Additional

6. Privacy of Trajectory Microdata Sharing and Release

This Section describes the privacy issues and proposed technologies to the protection of released or shared trajectory microdata. More details on the risks, attacks, and protection methods can be found in the survey by Fiore *et al.*¹¹⁴

6.1. Unique Characteristics

A trajectory microdata set is a microdata set that contains trajectory data. Thus, each row in the dataset represents an individual and each column represents an attribute of the individual. At least one of the attributes must correspond to a trajectory, that is, a list of spatiotemporal points. Additionally, other columns may correspond to other attributes of the individual.

Trajectory microdata sets are special because the location information included in them can be considered both as quasi-identifiers and sensitive information. Let us illustrate this using an example. The location information of a certain individual contains several recurring positions, some of them occur at night-time hours and some of them at daytime hours. Most likely, those positions correspond to the place of residence and the workplace of the individual, respectively. It is also very likely that the combination of place of residence and workplace is unique (although it may not be the case in all situations), and so the individual can be uniquely identified. This would allow an attacker to perform a record linkage attack if, for example, the attacker has background information on the workplace of their victim. If the location information also includes recurring visits to a medical institution, then we also have attribute disclosure, since the attacker can be confident that the individual has some chronic medical condition (if the medical institution is specialised in some kinds of conditions, then the information the attacker gets is more precise).

Trajectory microdata is prone to privacy attacks on individual users because of two defining characteristics. Trajectory data are highly unique and hard to anonymise. Experiments show that 50% of mobile subscribers in a database with 25 million different users can be uniquely detected with the knowledge of only 3 frequent locations. Additionally, experiments have shown that only very coarse generalisation (e.g., at the city level) produces reliable anonymised trajectory microdata sets.

In the following, we explore the risks and attacks that trajectory microdata are subjected to and methods proposed in the literature to anonymise those kinds of datasets. Note that we are interested

¹¹⁴ Fiore, M., Katsikouli, P., Zavou, E., Cunche, M., Fessant, F., Le Hello, D., Aivodji, U., Olivier, B., Quertier, T. and Stanica, R., 2020. Privacy in trajectory micro-data publishing: a survey. *Transactions on Data Privacy*, 13, pp.91-149.

in releasing or sharing the anonymised datasets with individual records, not aggregates or models obtained from the datasets.

6.2. Risks

6.2.1. Record linkage

One of the main risks in the sharing or release of a microdata set, be it a statistical database or a database with mobility data is that of record linkage. In a record linkage attack, an attacker attempts to uniquely match record in a published dataset to a record previously (partially) known by the attacker. The previous knowledge of the attacker might determine the kind of protection necessary to avoid such attacks.

In Section 4, we introduced the concepts of identifiers, quasi-identifiers, and confidential attributes in microdata sets. In the context of mobility data, we need to include a kind of attribute known as location-based quasi-identifiers, proposed first by Bettini *et al.*¹¹⁵. A location-based quasi-identifier (LBQID) is defined as “a spatiotemporal pattern specified by a sequence of spatiotemporal constraints each one defining an area and a time span, and by a recurrence formula”. In simple terms, LBQID are certain combination of spatiotemporal points that uniquely identify a record in a database or an individual.

An additional challenge to the protection of mobility data is that these kinds of data display high unicity and regularity. A high unicity in the data means that trajectories are quite unique between different individuals. A good example of this unicity is home locations: typically, very few people will share the same home address. Combining the home address with the work address further single out concrete individuals. As mentioned before, 50% of mobile subscribers in a database with 25 million different users can be uniquely detected with the knowledge of only 3 frequent locations. On the other hand, the trajectories followed by individuals tend to be highly regular¹¹⁶, that is, their mobility traces tend to follow a very stable pattern. For example, during workdays, most people will follow a trajectory from home to the workplace and back to home.

Next, we describe a series of record linkage attacks depending on the information known by the attacker.

A straightforward record linkage attack is possible when an attacker knows a subset of locations visited by her victim, or a *subtrajectory* of a full trajectory of a user. Due to the huge unicity of trajectory data, a small subset of positions can be enough to single out individuals. Methods to protect based

¹¹⁵ Bettini, C., Wang, X.S. and Jajodia, S., 2005, August. Protecting privacy against location-based personal identification. In Workshop on Secure Data Management (pp. 185-199). Springer, Berlin, Heidelberg.

¹¹⁶ Gonzalez, M.C., Hidalgo, C.A. and Barabasi, A.L., 2008. Understanding individual human mobility patterns. *nature*, 453(7196), pp.779-782.

on k -anonymity or similar, described in Section 6.3 as indistinguishability-based methods, try to make any subset of positions (any subtrajectory) be shared by at least k individuals, reducing the unicity of the data. It is also possible that the attacker has information on the trajectory of her victims, potentially with the same granularity as in the target database, but with different sampling times. For example, an attacker might obtain mobility information from geo-tagged photos in her victim's social media and intend to find her complete record in a target database (possibly to learn some confidential attribute of her victim). The information the attacker has might have the same granularity, but the sampling times are different. Ma *et al.*¹¹⁷ propose a series of techniques to define the similarity of noisy or differently sampled trajectories.

Instead of exploiting the unicity of mobility data, an attacker might exploit its regularity. De Mulder *et al.*¹¹⁸ propose a record linkage attack based on mobility models. The attacker knows a Markovian model of the pathing of her victim, that is, it knows a model that indicates the probabilities of moving from some position to another, irrespective of all previous or future positions. After accessing the released mobility dataset, the attacker can build a similar Markovian model for each of the individual records comparing them to his own. The authors report a success rate of 80% in a dataset with 100 individuals.

6.2.2. Attribute linkage

While record linkage attacks rely on the unicity of mobility microdata to single out individuals, attribute linkage attacks exploit the *homogeneity* of the data to discover sensitive information about the target individual. As an example, consider a mobility dataset that has been anonymised using some method based on k -anonymity. In this example, an adversary is not capable of linking her victim to an individual record in the released dataset since the background information of the attacker matches k or more different individuals. However, if the confidential attributes of the k individuals (which may also be some location) are highly homogeneous, then the attacker might infer the value for this attribute with high confidence. The method by Zhen *et al.*, presented in Section 6.3.2.5, leverages on the extensions of k -anonymity, namely ℓ -diversity and t -closeness, to reduce the homogeneity of confidential attributes in the k -anonymous groups.

6.2.3. Probabilistic attacks

A third type of attacks are probabilistic attacks. In this setting, an attacker can refine some (probabilistic) belief in the values of some confidential attribute from the user after having access to

¹¹⁷ Ma, C.Y., Yau, D.K., Yip, N.K. and Rao, N.S., 2010, September. Privacy vulnerability of published anonymous mobility traces. In Proceedings of the sixteenth annual international conference on Mobile computing and networking (pp. 185-196).

¹¹⁸ De Mulder, Y., Danezis, G., Batina, L. and Preneel, B., 2008, October. Identification via location-profiling in GSM networks. In Proceedings of the 7th ACM workshop on Privacy in the electronic society (pp. 23-32).

the dataset. These attacks are a generalisation of attribute linkage attacks and exploit homogeneity in the released database to improve the knowledge of the attacker by observing the dataset. Methods in Section 6.3.3 deal with protection mechanisms against these attacks.

6.3. Privacy methods

6.3.1. Disclosure risk mitigation

Mitigation strategies follow the utility-first anonymisation approach. These mechanisms do not provide any formal privacy guarantees but aim to reduce reidentification risks by applying different techniques, such as noise addition, generalisation and coarsening with heuristic parameter choice. After the application of such techniques, the disclosure risk is calculated (for some objective disclosure prevention, *i.e.*, identity disclosure or attribute disclosure). If the obtained risk is still too high, the techniques are applied with more strict parameters. Several of the following techniques can be applied both in the context of location-based services and in static trajectory microdata sets.

One of the first mechanisms introduced within the general anonymisation literature, and within the location privacy literature is noise addition or *obfuscation*. Agrawal and Srikant¹¹⁹ introduced a method for privacy-preserving data mining in which users introduce noise drawn from a uniform or Gaussian distribution to their sensitive attributes. Later, the data aggregator can reconstruct the original distribution from the noisy data to train a classifier. Note that techniques to enforce ϵ -differential privacy use noise addition, but differential privacy calibrates the added noise to achieve strict privacy guarantees.

Cloaking techniques aim at reducing the granularity of the location data, both temporal and spatial. Hoh *et al.*¹²⁰ reduce the temporal granularity to prevent home identification. First, they mount an attack to identify the home location from 239 mobility traces spanning one week and with sample frequency of 1 location per second while the vehicles are switched on. The authors report 85% of the homes are identified in this scenario. Then, 75% of the points in the trajectories are dropped (sample frequency is reduced to 1 location per 4 seconds), resulting in a home identification rate of 40%.

Song *et al.*¹²¹ proposed the *segmentation* of trajectories. The authors measure the risk of reidentification by a uniqueness metric, that is the fraction of individuals in the dataset which are uniquely identifiable by a set of spatiotemporal points. To reduce uniqueness, the trajectories of

¹¹⁹ Agrawal, R. and Srikant, R., 2000, May. Privacy-preserving data mining. In Proceedings of the 2000 ACM SIGMOD international conference on Management of data (pp. 439-450).

¹²⁰ Hoh, B., Gruteser, M., Xiong, H. and Alrabady, A., 2006. Enhancing security and privacy in traffic-monitoring systems. IEEE Pervasive Computing, 5(4), pp.38-46.

¹²¹ Song, Y., Dahlmeier, D. and Bressan, S., 2014, January. Not so unique in the crowd: a simple and effective algorithm for anonymizing location data. In PIR@ SIGIR.

users are split and assigned to new different identifiers, mimicking partially consistent protection mechanisms for location-based services. While the method, reduces the uniqueness of trajectories, but anonymised trajectories remain highly unique. These methods do not perturb any of the spatiotemporal information but may reduce the utility of the data sets when the intended analysis requires the mobility traces for users over long periods of time.

SwapMob¹²², as described in Section 6.3.2.4, can also be used to protect static trajectory microdata sets. In SwapMob, partial trajectories are swapped among users after they meet at some point.

6.3.2. Indistinguishability-based methods

This strategy aims at reducing or eliminating the uniqueness in the quasi-identifiers (in this case, the location-based quasi-identifiers) and typically are based on k -anonymity and its extensions. By applying techniques such as noise addition, generalisation, location coarsening, and micro-aggregation, these strategies try to produce groups of individuals for which their quasi-identifiers are the same, so that an adversary with some background information cannot distinguish between individuals in such groups.

6.3.2.1. GLOVE

Gramaglia and Fiore¹²³ study the difficulty of enforcing k -anonymity in trajectory databases and propose the GLOVE algorithm, based on a specialised generalisation technique.

The proposed algorithm aims to protect respondents against record linkage attacks and does not make any assumptions on the attackers' knowledge (attackers might know full trajectories of their victims). The authors' criterion towards indistinguishability of records is k -anonymity, and aim at anonymising full trajectories of respondents, using generalisation and (possibly) suppression of spatiotemporal points.

The authors define a metric for anonymisability of a given mobility trace named k -gap, which estimates how difficult it is to hide a given trace in a dataset, according to the accuracy loss required to make a mobility trace of respondent a indistinguishable from other $k - 1$ traces. This k -gap for respondent a is denoted as $\Delta_a^k \in [0,1]$. Assuming each sample spatial point is contained within a region of space (the authors consider regions of 100×100 metres) the *sample stretching effort* between two sample points indicates how much the regions containing these points must be stretched so that they are equal (a similar procedure is taken for the time dimension). This stretching

¹²² Salas, J., Megías, D. and Torra, V., 2018, September. SwapMob: Swapping trajectories for mobility anonymization. In International Conference on Privacy in Statistical Databases (pp. 331-346). Springer, Cham.

¹²³ Gramaglia, M. and Fiore, M., 2015, December. Hiding mobile traffic fingerprints with glove. In Proceedings of the 11th ACM Conference on Emerging Networking Experiments and Technologies (pp. 1-13).

is shown in Figure 8. Then, the k -gap between two trajectories is the average stretching efforts for each of the sample points contained in the two trajectories. Finally, the value Δ_a^k is obtained as the mean trajectory stretching efforts between trajectory a and the $k - 1$ closest trajectories (the closest trajectories are those with minimal *trajectory stretching efforts*). Studying the distribution of Δ_a^k for $k = 2$ in two example datasets show that most trajectories can be easily anonymised using generalisation, but a non-negligible number of trajectories remain unique and easily re-identifiable.




Figure 7: GLOVE Spatiotemporal generalisation procedure.

GLOVE takes as input a database with user trajectories and a parameter k , and works as follows. First, build a matrix that contains the stretch effort of every pair of trajectories. Then iterate until all trajectories have been anonymised. At each iteration, the two fingerprints that have not yet been k -anonymised and have the smallest stretch effort are merged by generalisation according to the stretching defined above. If the merged trajectory still does not contain k individual trajectories, it is reinserted in the database and its stretch effort to the rest of trajectories is computed. The anonymised database thus contains trajectories merged from k individual trajectories. A threshold on the k -gap is included so that trajectories that are too hard to anonymise (require too much generalisation) are suppressed from the database.

The generalisation required to obtain 2-anonymity is of 1 km and 1 hour with no suppression and of 0.5 km and 40 minutes with the suppression of 5% of the users in the database. The algorithm is quadratic in the number of records in the database and the average length of the trajectories.

A limitation of this approach is that the resulting anonymised data may be suitable to study the common mobility patterns of respondents or study on aggregates but are not suitable to conduct studies on outlying behaviours (outlying behaviours are highly unique by definition). This limitation is true for all mechanisms based on the indistinguishability of respondents.

6.3.2.2. KAM

Monreale *et al.*¹²⁴ propose an anonymisation method based on k -anonymity through the generalisation and suppression of spatial points (this approach does not consider time). In this case, the authors assume that an attacker may know a sub-trajectory (a trajectory contained in some other trajectory or trajectories) of her victim. In this work, full trajectories are considered private or sensitive information, while sub-trajectories are considered only quasi-identifiers and thus not harmful for privacy. The privacy guarantee offered by this method is as follows. Given an anonymity threshold $k > 1$, a trajectory microdata set D and an anonymised trajectory microdata set D^* , D^* is a k -anonymous version of D if and only if the subtrajectory known by the attacker either corresponds to no full trajectories in D or it corresponds to k or more full trajectories.

First, the generalisation method consists of extracting characteristic points of the trajectories in the original dataset, clustering these characteristic points, and obtaining the centroid of each of the clusters. The centroids are then used as generating points for a Voronoi tessellation of the area (using the Euclidean distance), and the trajectories are converted into visits to each of the cells, using the centroid as location. To ensure that each of the Voronoi cells contains at least k visits, less frequently visited adjacent cells are iteratively merged by recomputing the centroids of these areas and repeating the tessellation process.

Then the authors propose two anonymisation methods: KAM_CUT and KAM_REC.

In the KAM_CUT method, the generalised trajectories in the input dataset are used to build a prefix tree. Given an anonymity threshold k , the prefix tree is anonymised, *i.e.*, all the trajectories which are a subtrajectory of less than k full trajectories are pruned from the prefix tree. The anonymised prefix tree, as obtained in the previous step, is post-processed to generate the anonymised dataset of trajectories D^* .

The KAM_REC method works in a similar way but, instead of pruning all infrequent sub-trajectories, it tries to reinsert them back in the tree, by finding their longest subsequence of points that map to some popular subtrajectory either still in the tree or shared by at least k other trimmed sub-trajectories.

In terms of computational complexity, the generalisation algorithm is linear in the number of points in the dataset, KAM_CUT is also linear, and KAM_REC is quadratic in the number of points in the dataset.

¹²⁴ Monreale, A., Andrienko, G.L., Andrienko, N.V., Giannotti, F., Pedreschi, D., Rinzivillo, S. and Wrobel, S., 2010. Movement data anonymity through generalization. *Trans. Data Priv.*, 3(2), pp.91-121.

6.3.2.3. NWA/W4M

Abul *et al.*¹²⁵ introduces a trajectory database anonymisation method based on a relaxation of k -anonymity named (k, δ) -anonymity, whereby the indistinguishability requirement among k entries in the database required for k -anonymity is relaxed by some uncertainty δ . We first describe the concept of (k, δ) -anonymity.

An uncertain trajectory (depicted in Figure 9) is a cylindrical volume consisting of a circle of radius δ centered at each position (x_i, y_i, t_i) of the original trajectory. A possible motion curve (in red in the Figure 9) is any trajectory within the volume. Under (k, δ) -anonymity, any possible motion curve is considered indistinguishable. Two trajectories are co-located if both are possible motion curves of the other uncertain trajectory. A (k, δ) -anonymous set of trajectories is a set of trajectories of size at least k where all trajectories are co-located. A trajectory dataset is (k, δ) -anonymous if all trajectories in the dataset are included in a (k, δ) -anonymous set and the distortion between the original dataset and the anonymised dataset is minimised. Note that $(k, 0)$ -anonymity is the same as k -anonymity.




Figure 8: NWA uncertain trajectory and possible motion curve in red.

Next, the authors propose the Never Walk Alone (NWA) anonymisation algorithm.

¹²⁵ Abul, O., Bonchi, F. and Nanni, M., 2008, April. Never walk alone: Uncertainty for anonymity in moving objects databases. In 2008 IEEE 24th international conference on data engineering (pp. 376-385). Ieee.

NWA is developed along three main phases. First, NWA partitions the original trajectory database into equivalence classes with respect to the times, that is, form groups that contain all trajectories that have the same starting and ending points, up to some error caused by the sampling time (e.g., one-hour intervals). Next, perform micro-aggregation of trajectories, with the restriction that the clusters cannot be separated by more than δ . Finally, the clusters are released as (k, δ) -anonymous sets.

NWA can only deal with trajectories with equal, periodic sample times and has a quadratic time complexity. Trajectories anonymised with NWA can be distorted up to tens of kilometres.

Authors extend NWA to deal with trajectories with different sampling times. Wait for Me (W4M)¹²⁶ builds on NWA but substitutes the Euclidean distance by the Edit Distance on Real sequences and the Linear Spatiotemporal Distance. Using these distances, the first pre-processing step is no longer necessary, and the anonymised dataset suffers less distortion than when using NWA.

6.3.2.4. SwapLocations

Domingo-Ferrer *et al.*¹²⁷ propose trajectory database anonymisation mechanism, called SwapLocations, based on k -anonymity that does not introduce fake, perturbed, or generalised trajectories. The method is based on micro-aggregation of trajectories, but instead of substituting the trajectories of k -anonymous sets by the cluster representatives, locations are swapped among individuals in the k -anonymous sets. As in KAM, the privacy guarantee achieved by the proposed method ensures that an adversary with knowledge of some sub-trajectory of her victim cannot recover the whole trajectory of her victim with probability higher than $\frac{1}{k}$.

First, the authors propose a distance measure for trajectories which considers both the spatial and temporal dimensions of trajectories, is computable in polynomial time, and can cluster trajectories not defined over the same time span. The distance is computed in two steps: first, trajectories are linearly interpolated and sampled with an identical periodicity; second, the total Euclidean distance between contemporary points is computed. When two trajectories span different time intervals non-overlapping points are deleted, and the total distance is divided by the percentage of suppressed points.

Trajectories are clustered in groups of size k according to the proposed distance and using the microaggregation algorithm. Once the clusters are obtained, locations in the trajectories within the clusters are swapped among individuals, respecting some thresholds in space and time to limit the

¹²⁶ Abul, O., Bonchi, F. and Nanni, M., 2010. Anonymization of moving objects databases by clustering and perturbation. *Information systems*, 35(8), pp.884-910.

¹²⁷ Domingo-Ferrer, J. and Trujillo-Rasua, R., 2012. Microaggregation-and permutation-based anonymization of movement data. *Information Sciences*, 208, pp.55-80.

amount of distortion in the trajectories. If no switch is possible respecting the thresholds, the points are deleted.

While no points are introduced or distorted and the trajectories suffer small distortions, the proposed distance metric is not capable of handling trajectories that do not overlap, and therefore a large quantity of points is suppressed. The computation complexity of the proposed method is quadratic in the number of trajectories.

6.3.2.5. Zhen *et al.*

Zhen *et al.*¹²⁸ propose a trajectory dataset anonymisation mechanism that goes beyond k -Anonymity, introducing, ℓ -Diversity, and t -Closeness in their approach. As introduced in Section 4.2.3.3, k -anonymity can protect against record linkage and re-identification, however a k -anonymous dataset might still leak some confidential attributes of individuals. In the trajectory database context, consider an attacker with knowledge about a subtrajectory of her victim. If the trajectory dataset is k -anonymous, the attacker will not be able to distinguish her victim among the k individuals in the anonymous set. However, if all k trajectories contain the location of, for example, a medical institution, the attacker can be confident that her victim visits such medical institution. ℓ -diversity and t -closeness extend k -anonymity by ensuring a good representation of the confidential attributes (in this case of sensitive locations).

The method proposed by Zhen *et al.*¹²¹ follows a similar approach to GLOVE, representing the points of trajectories as the centroids of cells in the space and in some time intervals and with an additional input containing a list of Points of Interest (Pol). First, individual points are merged to achieve 2-anonymity by grouping the cells that contain such points, including any cell in-between. Additionally, other cells containing Pols are also merged (if necessary) to fulfil the ℓ -diversity and t -closeness criteria. Once points are merged and the 2-anonymity, ℓ -diversity and t -closeness criteria are fulfilled, the trajectories are further merged to achieve the desired k -anonymity.

The performance evaluation by the authors shows that the proposed solution can reduce Kullback-Leibler divergence by a factor of three while sacrificing an additional 30% of the spatial and temporal resolution over the baseline k -anonymity granted by GLOVE, with a similar quadratic time complexity.

6.3.3. Uninformativeness-based methods

Methods that follow this strategy are typically implemented with the use of ϵ -differential privacy or any of its relaxations. The objective of these methods is to reduce the information that an attacker

¹²⁸ Tu, Z., Zhao, K., Xu, F., Li, Y., Su, L. and Jin, D., 2018. Protecting Trajectory From Semantic Attack Considering k -Anonymity, ℓ -Diversity, and t -Closeness. *IEEE Transactions on Network and Service Management*, 16(1), pp.264-278.

would obtain on some attribute of an individual, or how their knowledge would change, if given access to the released dataset. Methods to enforce such privacy guarantees revolve around adding noise from some specific distribution and with parameters mandated by the privacy model.

6.3.3.1. Shao et al.

Shao *et al.*¹²⁹ propose a mechanism for the (ϵ, δ) -differentially private publication of individual trajectories. The goal of the proposal is to protect the true path of individual ships and use sampling and interpolation to achieve $(0, \delta)$ -differential privacy.

The authors propose two different approaches: Sampling First and Interpolation (SFI) and Interpolation First and Sampling (IFS), both achieving $(0, \delta)$ -differential privacy. In SFI, given a trajectory T , it is partitioned in groups of $k = \lceil 1/\delta \rceil$ points and a random point is chosen from each group. Then, the removed points are recovered by using cubic Bézier interpolation. On the other hand, in IFS the first step is to interpolate the curve in each time interval by using the cubic Bézier interpolation and then sample an alternative trajectory T_{mid} . Experimental work shows that SFI works better for smaller values of δ . We should note that most of the authors' choices for the value of δ in the presented experiments are not aligned with the restrictions set on δ by Dwork and Roth¹³⁰, that required δ to be negligible in the size of the dataset.

6.3.3.2. ϵ -DP Synthetic data generation

Chen *et al.*¹³¹ propose a mechanism for the generation of synthetic datasets of mobility data that satisfy ϵ -differential privacy.

This proposal uses an n -gram model, which describes the trajectories in the original dataset as transition probabilities based on the past $n - 1$ visited locations. This approach is based on a $(n - 1)$ -Markov model, where the probability that a user visits a location only depends on the past $n - 1$ visited locations, and not on the whole history of the user. To build such models, the authors build an exploration tree, starting with 1-grams and expanding to 2-grams, 3-grams and up to n_{max} -grams, where n_{max} is a parameter of the proposal. The n -grams have a count value associated to the number of instances of such sub-trajectories in the database, which are added Laplacian noise

¹²⁹ Shao, D., Jiang, K., Kister, T., Bressan, S. and Tan, K.L., 2013, August. Publishing trajectory with differential privacy: A priori vs. a posteriori sampling mechanisms. In International Conference on Database and Expert Systems Applications (pp. 357-365). Springer, Berlin, Heidelberg.

¹³⁰ Dwork, C. and Roth, A., 2014. The algorithmic foundations of differential privacy. Found. Trends Theor. Comput. Sci., 9(3-4), pp.211-407.

¹³¹ Chen, R., Acs, G. and Castelluccia, C., 2012, October. Differentially private sequential data publication via variable-length n -grams. In Proceedings of the 2012 ACM conference on Computer and communications security (pp. 638-649).

calibrated by the privacy budget ϵ . The obtained noisy n -grams can be published or used to generate synthetic trajectories.

An alternative approach called DP-WHERE by Mir *et al.*¹³² uses differentially private empirical cumulative distributions for home locations, work locations and commute distances around home locations, calculated from the original mobility dataset. These distributions are later used to sample home locations and a commute distance. This distance is used to draw a circle around the home location. Then a work location is sampled from the distribution, restricted by the previous circle.

6.3.3.3. $k^{\tau,\epsilon}$ -anonymity

Gramaglia *et al.*¹³³ extend k -anonymity to provide uninformative in the mobility data context. $k^{\tau,\epsilon}$ -anonymity requires any subtrajectory of size τ to be shared by at least k different full trajectories. Additionally, these full trajectories must not share more than $\tau + \epsilon$ points, *i.e.*, the trajectories must be diverse enough apart from the τ shared points. When τ equals the size of the full trajectories, this privacy model reduces to k -anonymity. This approach might seem similar to Zhen *et al.*¹²⁸, but while the latter imposes restrictions on the distributions of the non-shared points, the former just requires them to be different (except for ϵ points).

In order to enforce $k^{\tau,\epsilon}$ -anonymity, authors first present k -merge, an optimal low-complexity algorithm that generalises (sub-)trajectories with minimal loss of data granularity, and then build on it to propose $k\epsilon$ -hide, which enforces $k^{\tau,\epsilon}$ -anonymity on trajectory databases.

k -merge is an algorithm that merges k trajectories by iteratively generalising points in space and time so that each generalised point includes at least 1 point from each of the k original trajectories. The generalisations work in a similar way to that of GLOVE. The algorithm minimises the distortion caused by the generalisation and is linear in the number of spatiotemporal samples.

Authors use k -merge to achieve $k^{\tau,\epsilon}$ -anonymity for a single user in the dataset. First, time is discretised into intervals of size ϵ , named epochs. At the beginning of each epoch, $k - 1$ users (a hiding set) minimising the cost of k -merge are selected, and the corresponding samples are merged using k -merge. The authors extend this procedure to all users defining the $k\epsilon$ -hide, algorithm, which deals with choosing the $k - 1$ hiding sets for each user.

Performance evaluation shows that the method can attain $2^{\tau,\epsilon}$ -anonymity, with $\tau = \epsilon$. The solution retains a median accuracy of 1-2 km in space and less than 1 hour in time.

¹³² Mir, D.J., Isaacman, S., Cáceres, R., Martonosi, M. and Wright, R.N., 2013, October. Dp-where: Differentially private modeling of human mobility. In 2013 IEEE international conference on big data (pp. 580-588). IEEE.

¹³³ Gramaglia, M., Fiore, M., Tarable, A. and Banchs, A., 2017, May. Preserving mobile subscriber privacy in open datasets of spatiotemporal trajectories. In IEEE INFOCOM 2017-IEEE Conference on Computer Communications (pp. 1-9). IEEE.

7. Privacy of Aggregated Mobility Data

7.1. Unique Characteristics

As discussed in previous sections, the collection and release of mobility data, especially the release of data containing individual records or mobility traces, may lead to serious privacy issues, including identity or attribute disclosure. In order to mitigate these issues, data owners often rely on publishing only aggregated mobility data, such as the number of users covered by a cellular tower at a specific timestamp, which is believed to be sufficient for preserving users' privacy. An example of this is the pilot collaboration, in 2019, of major telecommunication companies in Spain with the Spanish National Institute of Statistics (Instituto Nacional de Estadística, INE) to publish mobile-based aggregated mobility data in Spain¹³⁴. Since the COVID-19 outbreak, INE has published such aggregated data in a weekly basis. The procedure to publish aggregated mobility data can be summarised as the following two steps: a) group the original mobility records of mobile users in time slots, b) compute and publish aggregated statistics of each time slot, for example, the number of users covered by each base station. However, this belief has been demonstrated to be false by some studies where authors have demonstrated that individual trajectories can be recovered from aggregated data, or that individuals can be singled out in aggregated data. The reasons for these risks are, again, the high unicity and regularity of user mobility data.

7.2. Risks

Aggregated mobility data are subjected to different risks than real-time location data as in LBS and of individual trajectory microdata. The risks we consider in the following are the recovery of trajectories from aggregated data, that is, extracting individual trajectories from aggregates, and membership inference attacks, that is, the decision of whether the data from a user which the attacker has some information about has been used to compute the aggregates. Both attacks can in turn lead to the attacks presented in previous sections. For example, recovering a trajectory from an individual in an aggregated dataset can be used to link it to a trajectory included in some trajectory microdata set.

¹³⁴ INE – Mobilty studies from mobile devices data.

https://www.ine.es/en/experimental/movilidad/experimental_em_en.htm

7.2.1. Recovery of trajectories

Tu *et al.*¹³⁵ present an attack on aggregated mobility data that can recover individual trajectories from published aggregated data. The authors leverage the high unicity and regularity of user mobility data to successfully recover unique user trajectories. The recovery of unique trajectories is a first step to conduct other kinds of attacks, such as re-identification attacks and record linkage attacks.

The proposed attack is based on a cost minimisation problem and leverages the differences of activity during night and day times. Between 89% and 95% of individuals remain static during the night, which facilitates the prediction of the position of individuals based on their previous position. During daytime, the authors leverage the average speed of individuals to predict the most probable next locations. Finally, the authors exploit the regularity of mobility data to match user trajectories throughout different days. The authors test the attack on both a dataset collected through a mobile application consisting of data from 15,000 users, and a dataset collected by a network operator in China containing data from 100,000 individual users and can recover up to 91% of unique trajectories.

Authors propose generalisation and perturbation mechanisms applied to the raw data to reduce the success of their proposed attack.

7.2.2. Membership inference

Zhang *et al.*¹³⁶ and Pyrgelis *et al.*¹³⁷ propose membership inference attacks on aggregated mobility data. A membership inference attack tries to infer whether an individual's data have been used to compute such aggregate. These types of attacks are especially harmful when the presence in a released dataset reveals some confidential attribute. For example, medical institutions are increasingly relying on wearable devices to follow the activities of patients that suffer concrete conditions. These wearable devices record mobility data along with some other attributes (e.g., blood pressure). Finding out that an individual's data have been used to compute an aggregate dataset is enough to infer that such individual suffers from a concrete medical condition.

The attacks leverage on prior information, such as the knowledge of some mobility patterns of the victim, the victim's colleagues, or the presence of the victim in past aggregates, and the published

¹³⁵ Tu, Z., Xu, F., Li, Y., Zhang, P. and Jin, D., 2018. A new privacy breach: User trajectory recovery from aggregated mobility data. *IEEE/ACM Transactions on Networking*, 26(3), pp.1446-1459.

¹³⁶ Zhang, G., Zhang, A. and Zhao, P., 2020. LocMIA: Membership Inference Attacks Against Aggregated Location Data. *IEEE Internet of Things Journal*, 7(12), pp.11778-11788.

¹³⁷ Pyrgelis, A., Troncoso, C. and De Cristofaro, E., 2017. Knock knock, who's there? Membership inference on aggregate location data. *arXiv preprint arXiv:1708.06145*.

aggregates and use machine learning algorithms to classify aggregate datasets according to the presence or not of the victim's data.

Again, Pyrgelis *et al.*¹³⁸ propose generalisation, perturbation, and suppression to limit the success of membership inference attacks.

7.3. Privacy methods

As acknowledged by Tu *et al.*¹³⁵ and Pyrgelis *et al.*¹³⁷ mechanisms based on generalisation, perturbation and suppression reduce the risks of trajectory reconstruction and membership inference attacks. Such methods have been proposed in Section 6.3, and thus it is advisable to use trajectory microdata anonymisation mechanisms prior to computing and releasing any aggregate mobility data. Among the methods already proposed, Chen *et al.* in Section 6.3.3.2 is directly applicable to aggregate data, although it does not preserve temporal information.

7.3.1. Zhili Chen *et al.*

Zhili Chen *et al.*¹³⁹ propose a suite of differentially private mechanisms for the release of aggregate mobility data, where the aggregates are visit counts at different times.

The first of their proposals consists of the application of the Laplace mechanism on the visit count histograms, that is, the addition of noise from a Laplacian distribution calibrated by the privacy budget ϵ . The second of the proposed mechanism sets a threshold by which histograms with little changes between time periods are not applied any noise, conserving the privacy budget. The two last mechanisms leverage the difference in mobility patterns at day and night times to further conserve the privacy budget and improve data utility. The authors provide experimental results for utility and resistance against the trajectory recovery attack by Tu *et al.*¹³⁵ on two datasets. Results suggest the provided mechanisms can thwart the trajectory recovery attack with a reduced data quality loss.

¹³⁸ Pyrgelis, A., Troncoso, C. and De Cristofaro, E., 2020. Measuring membership privacy on aggregate location time-series. *Proceedings of the ACM on Measurement and Analysis of Computing Systems*, 4(2), pp.1-28.

¹³⁹ Chen, Z., Kan, X., Zhang, S., Chen, L., Xu, Y. and Zhong, H., 2019. Differentially private aggregated mobility data publication using moving characteristics. *arXiv preprint arXiv:1908.03715*.

8. Conclusions

This document has analysed the regulations and risks to which data, and in particular mobility data is subject to, along with several techniques in the literature to protect mobility and location data. In the MobiDataLab project, we aim at developing a service and protocols to support and encourage the sharing of mobility data, and so methods aimed at the protection of trajectory microdata and aggregated data seem of a highest interest than those aimed at the protection of LBS, since data collection is likely to be performed by third parties before engaging in any sharing or release activities.

Regarding the protection methods, we acknowledge that as of today there is no standardised metric to compare their effectiveness, as there is no standard benchmark that would allow us to do so. The main reason for not having a benchmark is the lack of a standard or well-established dataset of mobility data in the scientific community. Currently, each publication uses its own dataset, and, in some cases, these are generated synthetically. Therefore, a comparison between the methods would be unfair even if the same measures are used. A secondary reason is that there is no standard privacy measurement technique applicable to all privacy methods, as most of them require specific techniques to realistically measure the risk of privacy leakage. Although several privacy models are well established, the methods to enforce such privacy models are often challenging to compare.

Regarding the utility of the resulting anonymised data, general metrics such as the spatiotemporal distortion or the mean squared error between original and anonymised datasets are often used, along with the number of suppressed locations. For specific uses, such as machine learning, we recommend training the models both on the original and the anonymised datasets to evaluate the loss of, for example, accuracy of the models.

9. Annexes

Table 4: Overview of strengths and weaknesses of techniques assessed by Article 29
Working Party in Opinion 5/2014¹⁴⁰

	Is singling out still a risk?	Is linkability still a risk?	Is inference still a risk?
Pseudonymisation	Yes	Yes	Yes
Noise addition	Yes	May not	May not
Substitution	Yes	Yes	May not
Aggregation or K-anonymity	No	Yes	Yes
ℓ-diversity	No	Yes	May not
Differential privacy	May not	May not	May not
Hashing / tokenisation	Yes	Yes	May not

¹⁴⁰ Article 29 Data Protection Working Party (2014). Opinion 05/2014 on anonymization techniques (829/14/EN WP216), p. 24.

| MobiDataLab consortium

The consortium of MobiDataLab consists of 10 partners with multidisciplinary and complementary competencies. This includes leading universities, networks, and industry sector specialists.



[@MobiDataLab](https://twitter.com/MobiDataLab)
[#MobiDataLab](https://twitter.com/MobiDataLab)



<https://www.linkedin.com/company/mobidatalab>

For further information please visit www.mobidatalab.eu



MobiDataLab is co-funded by the EU
under the H2020 Research and
Innovation Programme (grant
agreement No 101006879).

The content of this document reflects solely the views of its authors. The European Commission is not liable for any use that may be made of the information contained therein. The MobiDataLab consortium members shall have no liability for damages of any kind that may result from the use of these materials.