Location Privacy and Utility in Geo-social Networks: Survey and Research Challenges

Zohaib Riaz, Frank Dürr, Kurt Rothermel
Institute for Parallel and Distributed Systems
University of Stuttgart, Germany
Email: {zohaib.riaz, frank.duerr, kurt.rothermel}@ipvs.uni-stuttgart.de

***Abstract*—Location information sharing on popular online social networking platforms like Facebook and Foursquare brings mutual benefits for the users of these platforms (e.g., free location-based services) as well as the platform providers (e.g., location-based businesses). An obvious problem however that impedes these mutual benefits are privacy concerns related to location data of users, which also curb their active participation. In this paper, we analyze the role of existing location privacy-preserving mechanisms in minimizing this mutual loss of benefits. Our analysis reveals that most existing mechanisms either ignore social platform related user-privacy concerns or they disregard location data-quality related demands of the platform providers. Moreover, we also point out concrete research gaps and implementation issues related to existing privacy mechanisms.**

I. Introduction

Among other personal information, today's mobile devices sense the location of their users, e.g., using GPS, which enables a wide variety of location-based services (LBSs). A popular class of these LBSs are *geo*-social networks (GSNs) that range from people-locator services, e.g., Life360 for monitoring whereabouts of family members, to broader social networks, e.g., Facebook, Twitter, or Foursquare. It is no secret that today's GSNs, in general, can sustain free offering of their services to users by harvesting on the collected personal information including users' location data. Overall, the LBS market is projected to reach a $43.3 billion worth by 2019 based dominantly on location-based targeted advertising [1]. A core issue however that the location-based advertising industry faces today is low data quality, i.e, inaccurate location information, which is known to adversely effect most of the targeted ads (86% according to Telenav-Thinknear [2]) and yields low return-on-investments for marketers.

To improve this state of affairs, an attractive solution is to rely on user-reported accurate location information, i.e., *high-quality data*, such as that collected by existing GSN platforms. For instance, Foursquare uses first-hand location check-in information from its users to offer location intelligence to interested marketers [3]. However, to effectively mitigate the low location data-quality problem faced by the advertising industry, these GSN providers will need to source the larger share of the required location data, for example, by further encouraging the online information-sharing activity of users.

From GSN users' perspective, the fundamental hindrance in actively sharing their location information are the obvious privacy concerns. Recent studies show that people perceive their location data to be as sensitive, if not more, as their health data [4], [5]. Scientific works affirm that location data can indeed reveal private information, far beyond mere location coordinates, including a person's identity [6], [7], their daily pursuits [8], and their religious/political/sexual inclinations [9]. Thus users may naturally desire to control how and with whom their personal (location) information is shared (93% American users according to a 2015 Pew survey [10]).

Overall, the two parties, GSN users and platform providers, have somewhat conflicting goals, personal privacy and location data utility respectively, whose fulfillment is actively hindered in our view by their mutual non-cooperation. On the one hand, the GSN providers seem unwilling to take additional user-privacy enhancing measures since they may hinder collection of high-utility location data. On the other, the GSN users (or LBS users in general) react to their privacy concerns by reporting less data (as in switching off location tracking features in apps or checking-in less frequently [11]) as well as by reporting false location information [12].

In this paper, we survey location privacy-preserving mechanisms from existing literature and try to assess their role in addressing location privacy concerns of users and the data-quality concerns of GSN providers. To this end, we first identify the fundamental privacy demands of GSN users as well as the location-data related utility demands of GSN providers based on user studies as well as industrial reports. We term these demands surrounding location information as the *privacy-utility* requirements. We then categorize and analyze the different classes of existing location privacy mechanisms in view of their satisfaction of these privacy-utility requirements. In doing so, our analysis uncovers the high-level design limitations of the various classes of existing mechanisms, thus highlighting future research challenges. In particular, it reveals the fact that most existing location privacy mechanisms either do not address providers' utility concerns regarding location data or they disregard provider related user-privacy concerns. To summarize, we contribute by:

1) Specification of the privacy-utility requirements surrounding location information from the point of view of users, their social connections, and the GSN providers.
2) A classification and survey of existing location privacy mechanisms that are designed for use with GSNs.
3) Critical analysis of the existing classes of mechanisms with respect to the privacy-utility requirements, and as

978-1-5386-7493-2/18/$31.00 © 2018 IEEE
To appear in the proceedings of 16th Annual Conference on Privacy, Security and Trust (PST 2018), August 28-30, 2018, Belfast, Northern Ireland, United Kingdom.

a result, identification of future research challenges.

In our view, the above contributions may help industrial practitioners as well as the privacy research community to better understand the differences in their goals and find mutual grounds for the design of high privacy-utility systems. We also note that this work is significantly different from other surveys on location privacy research [13], [14] which focused on user-privacy concerns in use of LBSs in general and are not up-to-date with existing privacy-utility challenges in GSNs.

The rest of this paper is structured as follows. First, we define the GSN system model in Sect. II followed by the description of the privacy-utility requirements in Sect. III. In Sect. IV, we briefly survey the existing location privacy literature. The review of existing mechanisms with regards to the privacy-utility requirements is presented in V along with the identified research challenges. Finally, we conclude the paper in Sect. VI.

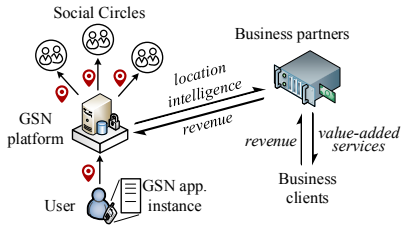

Fig. 1. The GSN eco-system with its four major entities: (1) Users, (2) Social Circles (SCs), (3) GSN Platform, and (4) Business Partners.

II. System Model

We represent the GSN eco-system by four main entities as shown in Fig. 1, namely, the users, their social circles, the GSN platforms, and third-parties who partner with the GSN providers in location-based businesses.

Users: The users of GSN platforms may be motivated to share their location information for several reasons such as reporting interesting events, earning discounts by checking-in to different venues, or in response to location requests from their social connections [15]. To perform location sharing, a user may access the platform's content sharing capabilities, such as check-ins or location tagging, through a browser-based or a dedicated mobile application on a location-enabled device such as a smartphone.

Social Circles: A typical user of GSN platforms may have a number of social connections which can generally be grouped into a few social circles (SC) such as family, friends, or colleagues, based on the nature of social interactions, level of closeness, and trust, etc. A number of studies [16], [17], [18] suggest that the comfort in sharing personal information (e.g., location) is highly dependent on the SCs who may view it. Hence users may share location data of varying granularity (precision) and quantity with different SCs.

GSN platforms: These platforms are operated by the GSN service providers, and connect users with their SCs. To this end, a platform facilitates its users in defining the members of their SCs and in sharing their location information with these SCs through location check-ins and other geo-tagged content. Moreover, these platforms typically allow users to express their privacy preferences as a set of *location-sharing rules* which control the times, locations, and the precision with which their location is made accessible to individual SCs [16], [18], [19].

Business Partners: To generate revenue from the gathered location information on users, the GSN providers, either themselves (e.g., Foursquare [3]) or in collaboration with third-party business partners (as done by Facebook [20]), offer value-added services to business clients. Without loss of generality, in this paper, we will limit these value-added services to the widely popular location-based targeted advertising and the collaborating third-party partners to marketing companies. For targeting customers, i.e., users of GSN platforms, these companies acquire "location intelligence" about them from the GSN providers (and other location publishers such as mobile network providers) and use it to run paid ad-campaigns for their clients businesses, who wish to increase customer foot-traffic to their stores. The quality of the acquired location intelligence, in terms of the accuracy and precision of its source location data, directly affects the return-on-investments for the business clients as a result of running ad-campaigns. Thus, GSN providers risk losing profitable collaborations with marketing partners if they are unable to collect sufficiently accurate location information on users.

III. Privacy-Utility Requirements

In the following discussion, we specify GSN users' privacy-utility requirements regarding their location information followed by the utility requirements of GSN providers. We base the specification of user related requirements on user-studies that revolve around their interaction with GSNs, e.g., [16], [18], [21], [17], [15], [22] as well as on the privacy-utility goals that have popularly been addressed in the location privacy literature. For specifying the utility requirements of GSN providers, we look at various industrial reports, e.g., [23], [24], [2], [25]. While we do not claim an exhaustive definition of these requirements, we have attempted to keep them sufficiently precise so as to highlight the differences between the various privacy-preserving mechanisms and to reveal interesting research gaps in our later discussion.

A. User Privacy Requirements

The social nature of interactions on GSN platforms may require from users to reveal their actual identities. Considering this, we limit the scope of privacy threats as well as requirements in this paper to those that go beyond user-identity information and span the additional personal information that may be inferred from location data of the users by potential recipients, i.e., the SCs and the GSN providers.

To illustrate the problem, Fig. 2 shows the "canvas" of historical location visits of a hypothetical user without timing information for simplicity. An attacker, for example in the role of a (malicious) "friend" belonging to an SC of the user, or the GSN provider itself, may have access to such information on

target users after extended interactions with them. By using online venue directories such as Foursquare or Yelp, the attacker may find out the types of places visited by users, i.e., the underlying location semantics, thereby drastically increasing the level of privacy threat. For instance, an individual visit of the user (single location update) with associated semantics of a hospital may reveal serious health issues. Extending semantic information to multiple locations, the attacker might infer movement habits and interests of the user, as shown in [8] and [9] repectively. Even without semantic information, location history information reveals the times users spent in different locations and the distance spanned by their movements, thus, making them vulnerable to identification [7] (even if they use pseudo-IDs) and personal security risks such as stalking [26] or even home robberies [27].

We will now discuss the privacy requirements of users in relation to their SCs and the GSN providers.

1) Privacy Requirements against SCs: The motivation for users to define privacy requirements regarding their SCs comes naturally from the fact that careless sharing of location information can result in regrets as well as social repercussions [15]. We classify these requirements into two types, namely, requirements regarding *individual* and *multiple* visits, which is also how they have been addressed by various privacy-preserving mechanisms, e.g., [28], [29], [30], [31].

Individual visits: As reported by Tang et al. [21], users may wish to control the precision and accuracy of their published location information from both, the geographic as well as the semantic perspective. In other words, they may wish to hide their precise geographic location or their activity. Despite the innate coupling between these two aspects, making one imprecise may not imply imprecision for the other. For example, a user may wish to inform his friends about his geographic location without revealing his current activity, e.g., eating in a restaurant. This requires publishing of a *manipulated* (for example coarsened) location which *must be large enough* to also subsume other venues than the restaurant in order to hide user-activity. Such a manipulated location supports another important requirement, namely, *plausible deniability*, as discussed in [21], [32], which gives users the freedom to present multiple justifications for their presence in the published location.

Multiple visits: Users may desire to control the projection of their personalities, or *personas* [16], [18], [33], that is conveyed to their SCs as a result of repeated sharing of location information. For example, a user may only share his visits related to social activities with his friends (social persona). Similarly, he may only share his accurate location with co-workers during working hours and no or only coarse location at other times (work persona).

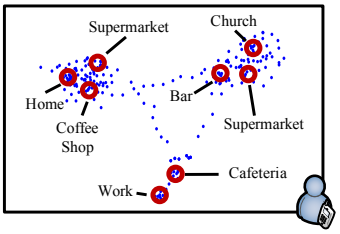

Fig. 2. Location history of a hypothetical user: location information may reveal the span of user movements and the kinds of places that they visit.

2) Privacy Requirements against GSN Providers: While legal regulations, such as the General Data Protection Regulation (GDPR) for the European Union, enhance the users' right of being forgotten[1] (data erasure), GSN platforms may be functionally restricted to temporarily store user data for service provision. For instance, as users may communicate their content asynchronously among each other, the GSN platform must store and manage it for possible deferred viewing by the recipient SCs. Similarly, prolonged retention of user data also enables popular platform features such as location history time-lines (e.g., by Google [34]). However, given frequent events of data breaches where popular companies lose sensitive user data [35], GSN users may justifiably fear for their privacy. Recent Pew surveys report that 69% and 76% of American users do not trust social media sites and online advertisers, respectively, to keep their data secure and private [10] and that 91% of users believe to have lost control over how their personal information is collected and managed by technology companies in general [22].

Thus, as privacy requirements, users may want to limit the total quantity of location information, i.e., their *location history*, that is aggregated by any individual GSN provider. In reference to Fig. 2 for instance, users may not wish the provider to know their complete canvas of their movements. Also, users may wish to hide details about their individual visits, e.g., to prevent unwanted ads. Thus the requirements of control over the geographic and semantic precision/accuracy of individual visits also apply against GSN providers.

B. Location Data-Utility Requirements

We now discuss the requirements on location data-utility that are imposed by the users, their SCs as well as the GSN providers.

1) Utility Requirements from Users: In return for sharing their location information, users may expect different *implicit* services, such as persona/impression management with their SCs, as well as various *explicit* LBSs from the GSN provider such as the listing of nearby friends or points-of-interests [36]. For explicit services, users may intend to share their location with the GSN provider exclusively while desiring a certain quality of the provided service (QoS) [37]. Meeting this QoS imposes a reciprocal requirement over the utility of the shared location information. For *individual* service requests, for example, the nearness of the points-of-interest returned by the GSN platform directly depends on the accuracy and precision of the reported location. More advanced services, such as personalized recommendations of previously unvisited venues [38], may require repeated sharing of location information (*multiple* visits) of a certain semantic type.

[1]Article 17 of the Regulation (EU) 2016/679 of the European Parliament and of the council of 27 April 2016.

(a) (b) (c) (d)

Fig. 3. Actual location of user (red) in (a) and its various representations (blue) after applying different location privacy preserving mechanisms in parts (b), (c), and (d). (b) *Spatial cloaking*: makes actual location imprecise, (c) *Location perturbation*: adds noise to actual location, (d) *Location dummies*: publishes true and a number of fake locations.

For the implicit services, the utility requirements on location information are mutually common between the users and the SCs and will therefore be discussed in the following section.

2) Utility Requirements from the SCs: In relations to the privacy requirements of the users against their SCs (as discussed in Sect. III-A1), we define the location-utility requirements imposed by the SCs of users as follows.

Individual visits: As described by Tang et al. in [21], location sharing with social connections may be purpose-driven or social-driven, which we argue, affects utility requirements of the SCs. In the former, location is shared on need basis with one or few persons, such as for coordination as in arranging meet-ups, as also confirmed by Cramer et al. [39]. Hence, in purpose-driven sharing, the quality of location information should be sufficiently high to fulfill the intended need. On the other hand, in social-driven sharing, location information is shared with a broader audience more for persona management than for meeting a particular need of a social connection. Hence, in contrast to purpose-driven sharing, users have relatively more choice over where, when, and how precisely to share their location.

Multiple visits: According to their social role, SCs may require from users to share their location information of a certain precision and regarding different semantic aspects of their movements, i.e., a certain persona (e.g., work, social, etc.). Note that this requirement coincides with users' desire to project certain personal impressions on their SCs.

Location representation: A further perceivable requirement imposed by SCs on the shared location information is its human-friendly representation for easy interpretation. This requirement becomes especially important if users choose to protect their privacy using existing mechanisms that may produce unnatural outputs. As shown in Fig. 3 (d), a popular mechanism, known as *location dummies* [40], generates multiple fake locations alongside the actual user location to confuse a non-trusted LBS provider. While this representation is suitable for querying an LBS provider for explicit services, it may be unsuitable for use in GSNs where multiple geo-coordinates for a single time instance may confuse the target SCs and additionally convey the fact that the user is employing a privacy-preserving mechanisms—an information that users may wish to conceal. In general, denial of location sharing requests from the SCs or sharing of overly imprecise/inaccurate location information has been shown to invite curiosity or convey feelings of distrust [41]. As seen in Fig. 3 (b) and 3 (c) respectively, spatial cloaking and perturbation mechanisms, on the other hand, generate protected locations with familiar representations that are also easily interpretable.

PRIVACY requirements of users

Caused by	Visit Type	Req.	Description
SCs & GSN Pr.	I	**P1**	*control geographic precision/accuracy*
		P2	*control semantic precision/accuracy*
SCs only	I	**P3**	*offer plausible deniability*
	M	**P4**	*restricted sharing (only persona)*
GSN Pr.	M	**P5**	*limit location history aggregation*

UTILITY requirements of users, their SCs, and the GSN provider

From	Visit Type	Req.	Description
Users	I/M	**U1**	*meet desired QoS (explicit LBSs)*
SCs	I	**U2**	*enable purpose-driven sharing*
	I/M	**U3**	*have interpretable representation*
Users & SCs	I/M	**U4**	*allow persona management (implicit services)*
GSN Pr.	-	**U5**	*be precise and accurate*
	-	**U6**	*constitute sufficient per-user location history (for audience profiling)*

Note: I=individual visit, M=multiple visits

TABLE I
USER-PRIVACY AND DATA-UTILITY REQUIREMENTS REGARDING LOCATION INFORMATION IN THE GSN ECOSYSTEM.

3) Utility Requirements from the GSN Providers: The location-data utility requirements of GSN providers, or any location data publisher, are driven by popular mobile advertising apps. According to recent market surveys [23], [24], [2], [25], these apps include proximity-targeting (targeting customers in the vicinity of client stores), audience targeting (targeting customers based on analysis of their location history), as well as attribution (measuring impact of location-based ads on client store foot-traffic). To enable these apps, the GSN providers ideally wish to:

1) *collect accurate location information* for proximity targeting or attribution etc.
2) *collect reasonably complete per-user location-histories* in order to infer personal traits such as their shopping habits (audience targeting).

The above discussed user privacy and utility requirements are summarized as **P1-P5** and **U1-U6** in Table I, respectively, and will be referred as such in the rest of this paper.

IV. STATE-OF-THE-ART: A BRIEF SURVEY

To survey the large body of existing location privacy mechanisms that are applicable to GSNs, we divide them into three broad categories, namely, *device-centric*, *(GSN) infrastructure-centric*, and *hybrid* mechanisms. The detailed taxonomy, which also dictates the order of discussion of the various mechanisms in this section, is shown in Fig. 4. While we try to discuss all representative works for each category of mechanisms, the following analysis may not exhaustively consider every mechanism given the page limit.

A. Device-centric mechanisms

Device-centric mechanisms assume that the user's device is the only component that may be trusted in the GSN-ecosystem.

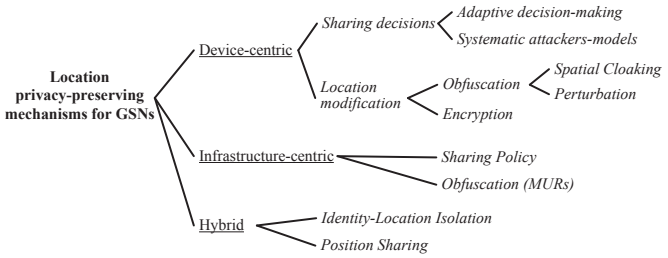

Fig. 4. The detailed taxonomy of location-privacy mechanisms for GSNs.

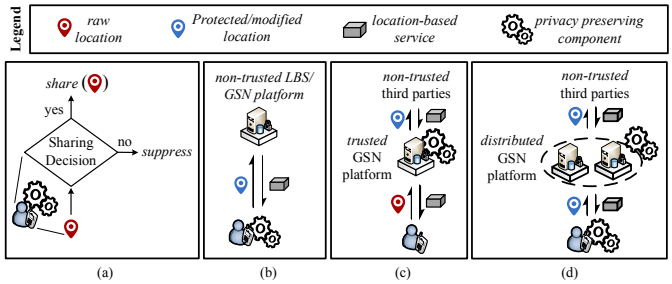

Fig. 5. Categories of location privacy-preserving mechanisms. (a) & (b) *Device-centric*. (c) *(GSN) Infrastructure-centric*. (d) *Hybrid* mechanisms. Privacy-preserving mechanisms runs on the user's device in (a) & (b), on the GSN platform in (c), and on both of these components in (d).

In other words, the user's location information must be protected before it leaves the device. Thus, the privacy algorithm executes on the device as a trusted software component and forms a gateway to location information for all requesting apps. For our following discussion, we define two subclasses of device-centric mechanisms, namely, the *sharing-decision based* and the *location-modification based* mechanisms as also depicted in Fig. 5 (a) and (b) respectively.

1) Sharing-decision based mechanisms: Sharing-decision based mechanisms determine *whether or not to publish the current user location*, ideally, without user intervention, i.e, in an autonomous fashion.

One form of these mechanisms treat users' manual sharing-decisions as the gold standard and try to learn from the users' example by applying machine-learning techniques. This methodology enables fine-grained *adaptive* decisions on individual location updates. In this regard, Bigwood et al. [42] demonstrated that decision learning over contextual features, such as time, the semantic place type, requesting SC, etc., can adaptively avoid privacy leaks that are conceded by, the relatively static, location-sharing rules (cf. Sect. IV-B). Other extensions add awareness of psychological factors (e.g., users' sharing tendency, trustworthiness of requester) [43]. The approach by Bilogrevic et al. [44] attempts to reduce user-burden of defining their privacy-preferences by using active learning and cost-sensitive classifiers. Another interesting work, Xie et al. [45], brings users in the sharing-decision loop while offering privacy recommendations, i.e., a few suitable sharing options to choose from to reduce the decision's complexity.

Since the above mentioned adaptive-decision mechanisms learn from the user's example, they at most only avoid privacy leaks as perceived by the users. In general, users may not fully perceive privacy threats resulting from information shared in the past, for example, due to its sheer volume and complexity as well as due to bounds on human rationality as suggested by Acquisti and Grossklags [46]. In this regard, Götz et al. [47] propose a mechanism to avoid privacy leaks, i.e., sharing of sensitive user context (e.g., being in a hospital), against a *systematic attacker* who knows the users' historically published context data. This mechanism models the attacker's knowledge of users' context switches as a Markov chain and uses suppression of the sensitive contexts for avoiding privacy leaks. The authors show that high correlation between sensitive and non-sensitive contexts, for example due to regular user movements as in home to work, forces their algorithm to additionally suppress some of the non-sensitive contexts to guarantee privacy for the sensitive ones.

2) Location-modification based mechanisms: These mechanisms form the second sub-class of device-centric mechanisms and protect user privacy by obfuscating or encrypting location information before its publication. As shown in Fig. 5 (b), these mechanisms may generically be used to preserve user privacy in GSNs as well as in other forms of LBSs.

Location obfuscation as defined in one of the earliest works by Dukham et al. [28] is the deliberate degradation of spatial location information quality by making it inaccurate, imprecise, or vague. Subsequently, most works have focused on adding imprecision by *spatial cloaking* (generalization as shown in Fig. 3 (b)) or by adding inaccuracy through *perturbation* (see Fig. 3 (c)).

Spatial Cloaking typically decreases the precision of location information by publishing a region containing the actual user location [29]. Since an attacker may attempt to prune this region by using map information (e.g., using locations of points-of-interests), other works attempt to hide user activity by generating semantically robust cloaking regions containing several semantically heterogeneous locations, e.g., bars, shopping malls, etc. [48], [49]. Instead of obfuscating every location update of users, mechanisms like that of Damiani et al. [50] offer to cloak only the sensitive semantic locations of users, e.g., a hospital, and allow users to configure their privacy preferences regarding these location types. However, in general, spatially cloaked locations are susceptible to pruning by *location-history-aware* attackers. For instance, in case of frequent location updates, the attacker may simply trim off those parts of an obfuscation region that are not reachable from their last reported user location within the elapsed time based on an estimate of the *maximum movement speed* of the user [37]. Avoiding such attacks needs special measures, such as intelligently delaying subsequent updates [30]. A more recent work by Riaz et al. [51] also highlights that mobility prediction attacks may also be possible based on sporadic location updates such as check-ins and can cause privacy breaches against semantic cloaking approaches.

For improved privacy guarantees, a recent class of *location perturbation* mechanisms assume location-history-aware attack models. To preserve location privacy under sporadic

location updates, Shokri et al. [31] propose to generate perturbed output locations that are probable candidates as per prior *spatial-location* distribution of the user. Other extensions [52], [53] explicitly model continuous movement as Markov chains over regions in space to appropriately perturb user location. Moreover, attempts have also been made to adapt the widely accepted notion of *differential privacy* [54] (from database systems) for location privacy of users [55], [53] as it originally offers *strong privacy guarantees against any background knowledge of the attacker*. More specifically, differential privacy allows privacy preserving querying of aggregate statistics over data from a user population such that an attacker is unable to infer any significant new information about individual users that was not already known. In contrast, the privacy requirements in the location-sharing setting must be met over data belonging to individual users rather than that of a population. Thus, so far, the adaptations of differential privacy for location privacy, as in location perturbation mechanisms [55], [53], cannot guarantee protection against *any background knowledge* of attackers. Notably, Xiao and Xiong [53] show that these adaptations must explicitly model attacks based on movement correlations, similar to the spatial cloaking mechanisms, when location updates are published frequently.

Location Encryption: Finally, a number of cryptographic protocols have been defined for delivering privacy-preserving *proximity notifications* among friends. While not trusting the GSN infrastructure, many of these mechanisms reasonably require honest execution of the privacy-preserving proximity protocol on the GSN platform over encrypted location information of users. A representative mechanism in this regard was presented by Mascetti et al. in [36]. Like other popular protocols, this mechanism overlays a rectangular grid over space to represent the space of user locations as a finite number of non-overlapping cells. The size of these cells is user-defined and enables control over the uncertainty with which their friends can locate them. Location information of users is shared with the GSN platform as encrypted index of the cell containing their actual location. Based on shared secrets between users and their friends, such as encryption keys and grid size information, this mechanism proposes protocols for the determination of (a) proximity with friends; (b) proximity along with cell-level location of friends. Baden et al. [33] propose a social network, called *Persona*, which allows fine-grained sharing of data between users by employing Attribute-based Encryption (ABE). While users may target information more easily using ABE to specific social connections who satisfy certain access attributes, the storage provider does not learn any user-information as it is managed in encrypted form.

B. Infrastructure-centric mechanisms

Infrastructure-centric mechanisms, as shown in Fig. 5 (c), assume that the GSN platform is trusted by its users. Hence privacy concerns only come from the users' SCs. Since the SCs typically know the real identities of users, our discussion here leaves out a number of infrastructure-centric mechanisms that aim at preserving users' anonymity in using LBSs (see [14] for a survey).

As a first line of defense for their privacy, GSN platforms allow users to define their *location sharing policy* with their SCs as a set of location sharing-rules [16], [18], [19], for instance, *share my location with my friends only after 5 pm*. In defining these rules however, users are known to face two major problems. First, as per the study by Benisch et al. [56], the initial definition of these rules is cumbersome for users as a large number of rules may need to be specified to fully express user-privacy preferences. Secondly, initial rule definitions are also often inaccurate [57], [58] and can result in privacy leaks. To address the first problem, Ravichandran et al. [59] have investigated the possibility of finding representative sets of sharing rules that may provide a basic template for users to adapt to their specific needs, thus, reducing user burden. Toch et al. [17] also suggest using properties, such as *location entropy* (in terms of the diversity of visiting users and their visiting intensities), to determine the default sensitivities of different semantic locations (e.g., home, work, etc.) which acts as a heuristic to define default sharing rules. To address the second problem of inaccurate initial rules, Sadeh et al. [60] conducted a user study to show that location-sharing rules are intermittently refined by users for increased accuracy compared to their initial definitions, and that automatic refinement by applying machine learning techniques outperforms manual refinement in terms of the number of incurred privacy leaks. Extended works integrate users into the rule refinement loop in order to enable more user-control over the refined sharing rules [61].

Beyond the definition of location sharing policies, few works have adopted the infrastructure-centric model for privacy-preservation in GSNs. Among these, the work by Ferni et al. [62] is arguably the most prominent. This work addresses the problem of privacy-preserving publishing of location-tagged resources, such as photos, in GSNs. Doing so however is not straight-forward as location-tagged resources may reveal co-location of users, e.g., in group photos. Hence determining a privacy preserving location tag requires consideration of privacy preferences of all users associated with the resource. Furthermore, since this mechanism also aims to secure the location and *absence* privacy of users, it allows users to specify their privacy preferences as spatio-temporal uncertainty about their presence or absence in different regions of their movement. The GSN platform then enforces these preferences by publishing *minimum uncertainty regions* (MURs) that result from the spatial generalization of users' location and temporal obfuscation (delaying) in resource publication. In doing so, this mechanism also ensures that MURs published nearby in time are not vulnerable to attacks based on knowledge of users' maximum movement velocity.

Other works relying on infrastructure-centric mechanisms aim at *privacy-preserving publishing* of location information. For example, Riboni et al. [38] propose a mechanism for GSN platforms to publish check-in statistics for context-aware location recommendations for mobile users. To this end, the

GSN platform uses differential privacy to perturb the check-in counts of various venues before reporting them to their business partners for location recommendations.

C. Hybrid mechanisms

As the name suggests, hybrid mechanisms rely on trusted user device as well as non-trusted GSN infrastructure for ensuring user privacy (see Fig. 5 (d)). To achieve this, these mechanisms typically decouple the storage of user information from GSN application functionality, and implement the storage part in a distributed fashion on multiple (instead of one) non-colluding storage providers. Hence no single component/provider constituting the GSN platform is able to access all user information. In this regard, we categorize existing mechansims into two types: those that isolate user location and identity information, and those that distribute users' location information among different components (Position Sharing mechanisms).

1) Identity-Location Isolation: The idea of a decentralized matching service to enable privacy-preserving LBSs has been similarly pursued in [63] and [64]. Jaiswal et al. [63] propose an mechanism that aims to avoid the aggregation of location information of users (as known to the mobile network provider in their system) and their interests (represented by users' queries to LBS providers) by any one entity. To achieve this, the network provider periodically encodes the user-reported locations as pseudo-locations (PLs) on a spatial grid. Similarly, the LBS providers generate pseudo-identifiers (PIs) for the venues in their database or for their registered users. These PLs and PIs are then shared with a third component, called the matching service, which is then able to answers user queries for nearby points-of-interests (POIs), without knowing the actual user identities or their locations.

Another extension by Guha et al. [65] proposes a cloud-based two-party implementation of the matching service called Koi. One of the cloud components in Koi, called the *matcher*, stores the identity information of entities (users or POIs) as well as their attribute values (e.g., location) without knowing their relationship. For example, the matcher may know that there is a user "Bob" and there is a location-attribute of "New York", but it does not know that this attribute belongs to Bob. The relationship between identities and attributes of users and POIs are stored with the second non-colluding component, from a different provider, called the *combiner*. By running a privacy-preserving protocol on pseudo-ids of the entities and their attributes, the matcher and the combiner implement a matching service to enable various LBSs without learning about each others data. However, in absence of enough cover traffic between the two components, the matcher can perform a traffic analysis attack to correlate subsequent matches between pseudo-ids of various entities and, thus, understand and uncover their relationship.

2) Position Sharing: The initial *Position Sharing* approach proposed by Dürr et al. [66] aims to avoid users' privacy concerns regarding insecure storage of their individual locations in a non-trusted GSN infrastructure. To this end, their proposed mechanism *divides* location information into a number of *imprecise location-shares* and *distributes* these shares, one each, to a number of non-trusted location servers that are operated by different providers. While these servers together form the storage infrastructure for location information, none of them individually holds the complete location information about the users. Hence if one or few of these location servers are compromised by an attacker, user *privacy degrades gracefully*. On top of such a distributed storage infrastructure, different LBS apps, including GSN apps, function by accessing location-shares from a limited number of these location servers and subsequently *fuse* these shares together to recreate location information of a certain precision (the more the shares, the higher the precision). By controlling the number of servers/shares that an application is authorized to access, users can control how precisely these apps, and consequently the SCs that they represent, view user location information.

The *divide and distribute* idea of the position-sharing mechanisms has been extended to protect single user trajectories [67] as well as inference-prone private information in location histories, such as frequently visited locations [9]. These extensions fundamentally change the definition of a *location share* stored at an individual location server to achieve their respective goals. For example, for protecting highly frequented semantic locations, Riaz et al. [9] defined shares in an online fashion by distributing user check-ins, as per their associated semantic information, to different servers such that individual servers do not learn the users' frequent locations.

V. Fulfillment of Privacy-Utility Requirements and Future Research Gaps

In Table II, we *approximately* summarize our upcoming analysis of the state-of-the-art location privacy mechanisms (per row) in terms of their fulfillment of the privacy-utility requirements (columns). We later also identify concrete research challenges as an outcome of this analysis.

A. How well are the requirements met?

We begin with the privacy requirements of *control over geographic and semantic precision/accuracy* of location data (**P1** and **P2** respectively). For **P1**, only *perturbation* mechanisms offer strong privacy guarantees by also considering location-history aware attackers. Moreover, only *spatial cloaking mechanisms* address **P2** (with semantic cloaking mechanisms) and that too *partially* by overlooking location-history based attacks. *Encryption*, *MURs*, and *location-identity isolation* mechanisms also ignore location-history based attacks to partially meet **P1**[2]. *Sharing decision-based* mechanisms, *sharing-rules* and *position-sharing* mechanisms (for multiple visits) do not meet requirements **P1** and **P2** as they generally publish accurate location information.

Conversely, all of these mechanisms can meet the utility requirements for explicit LBS (**U1**) and purpose-driven sharing

[2]Note that for Encryption mechanisms also, SCs who are authorized to decrypt individual locations can view current and past locations, and can exploit their spatio-temporal correlation for privacy attacks.

Privacy mechanism	Fulfillment of Privacy-Utility reqs.: ● full, ◒ partial, ○ unaddressed									
	P1	P2	P3	P4	P5	U1	U2	U3	U4	U5	U6
Device-centric:											
Sharing-decisions	○*	○*	○*	●	○	●	●	●	●	●	●
Spatial Cloaking	◒	◒	●	○†	○	●	●	●	○†	◒	●
Perturbation	●	○	○	○†	○	●	●	●	○†	◒	●
Encryption	◒	○*	●	○†	●	●	●	●	○†	○	○
Infrastr.-centric:											
Sharing-rules	○*	○*	○*	●	○	●	●	●	●	●	●
MURs	◒	○*	●	○†	○	●	●	●	○†	◒	●
Hybrid:											
Location-Identity Isolation	◒	○*	●	○†	○	●	●	●	○†	●	●
Position-Sharing	○*	○*	○*	●	●	●	●	●	●	●	◒

*: can improve using *Spatial Cloaking* or *Perturbation*.
†: can improve using *Sharing-decisions* or *Sharing-rules*.

TABLE II
FULFILLMENT OF THE PRIVACY-UTILITY REQUIREMENTS (CF. TABLE I) BY STATE-OF-THE-ART LOCATION PRIVACY MECHANISMS.

(**U2**) by either publishing accurate location information or by trading the level of obfuscation for QoS. Moreover, they also offer interpretable representations for location (thus meeting **U3**). From the GSN providers' perspective, all of these mechanisms, except the ones that obfuscate location (spatial cloaking, perturbation, and MURs), allow the *aggregation of accurate location data* at their servers to meet **U5**. To judge whether these mechanisms offer plausible deniability or not (requirement **P3**), we refer the reader to Fig. 3. It is reasonable to say that all those mechanisms whose output location is interpretable as a single/unique location (e.g., perturbation or sharing-rules/decisions) do not offer plausible deniability. As regarding the strict management of personas (requirements **P4** and **U4**), *sharing decision/rule-based* mechanisms as well as *position-sharing* mechanisms fulfill this requirement by design[3]. The remaining mechanisms focus more on protecting individual location updates and thus do not address requirements **P4** and **U4**.

Finally, privacy-threatening *aggregation of location-history information with the GSN provider* (property **P5**) is not prevented by existing mechanisms with the exception of *location encryption* and *Position Sharing*. Encryption mechanisms, however, naturally disable useful processing of location data thus disabling GSN providers' business models (unfulfilled **U5-U6**). Note that *Identity-location isolation* based mechanisms also do not fulfill **P5** by storing all location data of users, albeit lacking identity information, with a single provider in the storage infrastructure. Well-known works, e.g, by Gruteser et al. [6], have shown that location samples can be linked to re-create user trajectories (if published frequently) allowing recreation of all user movements (location histories), which can subsequently also reveal users' actual identities [7].

Overall, we observe that no single privacy mechanism addresses all privacy-utility requirements completely. Nevertheless, it seems viable to meet more requirements by integrating mechanisms that focus on protecting individual user visits with those focusing on multiple visits. For example, sharing-decision based mechanisms can be improved in conjunction with spatial cloaking to positively address properties **P1, P2, P3**. We have marked such possibilities of meeting a certain requirement in conjunction with *spatial cloaking* or *location perturbation* by the "*" symbol alongside the circles in the Table II. Similarly, the possibility of integration with *sharing-decisions* or *sharing-rule* based mechanisms to meet persona management requirements (**P4** and **U4**) is indicated by "†". Even after such integration of privacy-capabilities among these different mechanisms, notable challenges are still evident.

[3]Recall that position-sharing mechanisms allow users to authorize the various SCs about which location shares they can access.

B. Research Challenges

At first, existing mechanisms do not avoid location-history attacks for *semantic* cloaking approaches (see column **P2** of Table II). In GSNs, the provider and the individual social connections may possess varying knowledge of the user's location history and therefore form (potential) attackers of varying strength. In such a multi-attacker scenario, publishing an obfuscated visit is not straightforward even when employing a location-history-*aware* protection mechanism. As also noted in other works [47], [53], assuming an improper attacker strength for obfuscating locations may cause privacy breaches. Figure 6 shows an example in this regard. Here, if a strong attacker knowing the user's habit of visiting the church is assumed, the privacy algorithm generates a cloaking region (CR) CR_{strong} which contains the church along with other locations (restaurant in this case). In contrast, assuming a weak attacker who relies more on the general population behavior may result in CR_{weak}, which includes those locations that are popular at the time of visit. Now if CR_{strong} is published by the privacy algorithm, the weak attacker learns that "church" is a preferred place of visit for this user because it is included in CR_{strong}. On the other hand, if the weak attacker is assumed and, correspondingly, CR_{weak} is published, the strong attacker, who knows that the user prefers to go to church and the restaurant at this time, learns that the user must have visited the restaurant because it is part of CR_{weak}. Current works on semantic location obfuscation lack the handling of this and similar cases.

Secondly, it can be concluded from the above analysis that few approaches directly limit aggregation of location-history information with the GSN providers while still supporting the LBS business model (requirements **P5, U5, U6**). Although Position-sharing approaches present a promising solution to address this problem, it may be worthwhile to explore these approaches further. For example, it may be interesting to study

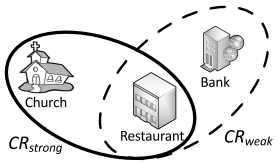

Fig. 6. CRs for *strong attacker* (solid line) and *weak attacker* (dotted Line).

the design of intelligent location history partitioning schemes that allow users to control the kind of audience profiling a provider may do about them (based on its portion of users' data) so as to, for instance, get relevant ads.

Thirdly, while Table II summarizes the algorithmic capabilities of existing mechanisms, it does not necessarily imply their practicality. For example, the performance aspect of mobile device-centric mechanisms, in general, has not been thoroughly evaluated. Especially with more complex mechanisms, e.g., those avoiding location-history based attacks, concrete (and perhaps open-source) implementations are needed to verify their viability on todays' mobile devices. Similarly, hybrid mechanisms also face implementation issues, such as the availability of enough location servers from independent providers to realize position-sharing mechanisms. Exploring further architectural designs for hybrid mechanisms that overcome these problems also seems promising.

VI. Conclusion

While offering popular platforms for social interactions, today's geo-social networks (GSNs) also collect location information of their users for location-based business gains. In doing so however, these providers incur user distrust and thus inhibit their active participation, especially in terms of their willingness to share personal location data. In this paper, we have assessed how the location-data related privacy concerns of GSN users may be addressed while still supporting the location-based business models of GSN providers. In this regard, we have critically reviewed existing scientific literature on location privacy-preservation in GSNs. As a result, we have provided a detailed assessment of the various classes of location privacy mechanisms in terms of their satisfaction of user-privacy requirements as well as the GSN providers' utility requirements regarding location data. Moreover, we have also pointed out open challenges for future research.

Acknowledgments

This work is a part of project *PriLoc* (Privacy-aware Location Management) of the University of Stuttgart, funded by the German Research Foundation (DFG) grant RO 1086/15-2.

References

[1] S. Smith, "Location based services market to reach $43.3 bn by 2019, driven by context aware mobile services - juniper research." https://tinyurl.com/y9pp9az3. Accessed: 2018-04-16. Published: 13th August, 2014.

[2] Telenav-Thinknear, "Location score index—mobile advertising's guide to location accuracy," Q2 2016. http://www.thinknear.com/library/location-score-index-q2-2016/. Accessed: 2018-04-16.

[3] Foursquare, "Foursquare location intelligence." https://enterprise.foursquare.com/. Accessed: 2018-04-16.

[4] Pew Research Center, "Public Perceptions of Privacy and Security in the Post-Snowden Era," 2014. https://tinyurl.com/yc3hfptu/. Accessed: 2018-04-16.

[5] D. Leibenger, F. Möllers, A. Petrlic, R. Petrlic, and C. Sorge, "Privacy challenges in the quantified self movement - an eu perspective," *Proceedings on Privacy Enhancing Technologies*, vol. 2016, no. 4, 2016. Conf. Presentation at PETS 2016.

[6] M. Gruteser and B. Hoh, "On the anonymity of periodic location samples," in *Proc. of the Second Int. Conf. on Security in Pervasive Computing*, SPC'05, (Berlin, Heidelberg), pp. 179–192, Springer-Verlag, 2005.

[7] J. Krumm, *Inference Attacks on Location Tracks*, pp. 127–143. Berlin, Heidelberg: Springer Berlin Heidelberg, 2007.

[8] N. Eagle and A. S. Pentland, "Eigenbehaviors: identifying structure in routine," *Behavioral Ecology and Sociobiology*, vol. 63, no. 7, pp. 1057–1066, 2009.

[9] Z. Riaz, F. Dürr, and K. Rothermel, "On the privacy of frequently visited user locations," in *Proceedings of the 17th IEEE Int. Conf. on Mobile Data Management (MDM)*, vol. 1, (Porto, Portugal), pp. 282–291, June 2016.

[10] Pew Research Center, "Americans' Attitudes About Privacy, Security and Surveillance," 2015. https://tinyurl.com/kvgfojh. Accessed: 2018-04-16.

[11] Pew Research Center, "Location Based Services," 2013. https://tinyurl.com/ycyq7gen. Accessed: 2018-04-16.

[12] Pew Research Center, "Teens, Social Media, and Privacy," 2013. https://tinyurl.com/m57j48h. Accessed: 2018-04-16.

[13] C.-Y. Chow and M. F. Mokbel, "Trajectory privacy in location-based services and data publication," *ACM Sigkdd Explorations Newsletter*, vol. 13, no. 1, pp. 19–29, 2011.

[14] M. Wernke, P. Skvortsov, F. Dürr, and K. Rothermel, "A classification of location privacy attacks and approaches," *Personal and Ubiquitous Computing*, vol. 18, no. 1, pp. 163–175, 2014.

[15] S. Patil, G. Norcie, A. Kapadia, and A. J. Lee, "Reasons, rewards, regrets: Privacy considerations in location sharing as an interactive practice," in *Proc. of the 8th Symposium on Usable Privacy and Security (SOUPS '12)*, (NY, USA), pp. 5:1–5:15, ACM, 2012.

[16] S. Patil and J. Lai, "Who gets to know what when: Configuring privacy permissions in an awareness application," in *Proc. of the SIGCHI Conf. on Human Factors in Computing Systems*, CHI '05, (NY, USA), pp. 101–110, ACM, 2005.

[17] E. Toch, J. Cranshaw, P. H. Drielsma, J. Y. Tsai, P. G. Kelley, J. Springfield, L. Cranor, J. Hong, and N. Sadeh, "Empirical models of privacy in location sharing," in *Proc. of the 12th ACM Int. Conf. on Ubiquitous Computing*, UbiComp '10, (NY, USA), pp. 129–138, ACM, 2010.

[18] J. Tsai, P. Kelley, L. Cranor, and N. Sadeh, "Location-sharing technologies: Privacy risks and controls," *ISJLP*, vol. 6, pp. 119–317, 2010.

[19] E. Toch, J. Cranshaw, P. Hankes-Drielsma, J. Springfield, P. G. Kelley, L. Cranor, J. Hong, and N. Sadeh, "Locaccino: A privacy-centric location sharing application," in *Proc. of the 12th ACM Int. Conf. Adjunct Papers on Ubiquitous Computing - Adjunct*, UbiComp '10 Adjunct, (NY, USA), pp. 381–382, ACM, 2010.

[20] Facebook, "Facebook marketing partners." https://facebookmarketingpartners.com/. Accessed: 2018-04-16.

[21] K. P. Tang, J. Lin, J. I. Hong, D. P. Siewiorek, and N. Sadeh, "Rethinking location sharing: Exploring the implications of social-driven vs. purpose-driven location sharing," in *Proc. of the 12th ACM Int. Conf. on Ubiquitous Computing*, UbiComp '10, (NY, USA), pp. 85–94, ACM, 2010.

[22] Pew Research Center, "The state of privacy in America — Pew Research Center," 2016. https://tinyurl.com/jfhhqfu. Accessed: 2018-04-16.

[23] Mobile Marketing Association, "Demystifying location data accuracy," 2015. https://tinyurl.com/y789jnr6. Accessed: 2018-04-16.

[24] Mobile Marketing Association, "How marketers are using location data and the road ahead," 2016. https://tinyurl.com/ycvhew6j. Accessed: 2018-04-16.

[25] IAB, "Iab mobile location data guide for publishers," 2016. https://tinyurl.com/y954tnk9. Accessed: 2018-04-16.

[26] The Telegraph, "Police warn over location tracking site following case of alleged Foursquare stalker," 2011. https://tinyurl.com/2ddd8yd. Accessed: 2018-04-16.

[27] "Please Rob Me." http://pleaserobme.com. Accessed: 2018-04-16.

[28] M. Duckham and L. Kulik, "A formal model of obfuscation and negotiation for location privacy," in *Pervasive Computing*, vol. 3468 of *Lecture Notes in Computer Science*, pp. 152–170, Springer Berlin Heidelberg, 2005.

[29] C. A. Ardagna, M. Cremonini, E. Damiani, S. De Capitani di Vimercati, and P. Samarati, *Location Privacy Protection Through Obfuscation-Based Techniques*, pp. 47–60. Berlin, Heidelberg: Springer Berlin Heidelberg, 2007.

[30] G. Ghinita, M. L. Damiani, C. Silvestri, and E. Bertino, "Preventing velocity-based linkage attacks in location-aware applications," in *Proc. of the 17th ACM SIGSPATIAL Int. Conf. on Advances in Geographic Information Systems*, GIS '09, (NY, USA), pp. 246–255, ACM, 2009.

[31] R. Shokri, G. Theodorakopoulos, C. Troncoso, J.-P. Hubaux, and J.-Y. Le Boudec, "Protecting location privacy: optimal strategy against localization attacks," in *Proc. of the ACM Conf. on Computer and communications security (CCS)*, pp. 617–627, ACM, 2012.

[32] J. I. Hong and J. A. Landay, "An architecture for privacy-sensitive ubiquitous computing," in *Proc. of the 2Nd Int. Conf. on Mobile Systems, Applications, and Services*, MobiSys '04, (NY, USA), pp. 177–189, ACM, 2004.

[33] R. Baden, A. Bender, N. Spring, B. Bhattacharjee, and D. Starin, "Persona: an online social network with user-defined privacy," in *ACM SIGCOMM Computer Communication Review*, vol. 39, pp. 135–146, ACM, 2009.

[34] Google, "Google - timeline." https://www.google.com/maps/timeline. Accessed: 2018-04-16.

[35] Information is Beautiful, "World's biggest data breaches & hacks." https://tinyurl.com/lgyx9lc. Accessed: 2018-04-16.

[36] S. Mascetti, D. Freni, C. Bettini, X. S. Wang, and S. Jajodia, "Privacy in geo-social networks: Proximity notification with untrusted service providers and curious buddies," *The VLDB Journal*, vol. 20, pp. 541–566, Aug. 2011.

[37] R. Cheng, Y. Zhang, E. Bertino, and S. Prabhakar, "Preserving user location privacy in mobile data management infrastructures," in *Privacy Enhancing Technologies*, vol. 4258 of *Lecture Notes in Computer Science*, pp. 393–412, Springer Berlin Heidelberg, 2006.

[38] D. Riboni and C. Bettini, "Differentially-private release of check-in data for venue recommendation," in *Proc. of PerCom*, pp. 190–198, March 2014.

[39] H. Cramer, M. Rost, and L. E. Holmquist, "Performing a check-in: Emerging practices, norms and 'conflicts' in location-sharing using foursquare," in *Proc. of the 13th Int. Conf. on Human Computer Interaction with Mobile Devices and Services (MobileHCI)*, pp. 57–66, ACM, 2011.

[40] H. Kido, Y. Yanagisawa, and T. Satoh, "An anonymous communication technique using dummies for location-based services," in *Proc. of the Int. Conf. on Pervasive Services (ICPS)*, pp. 88–97, 2005.

[41] D. Wagner, M. Lopez, A. Doria, I. Pavlyshak, V. Kostakos, I. Oakley, and T. Spiliotopoulos, "Hide and seek: location sharing practices with social media," in *Int. Conf. on Human-Computer Interaction with Mobile Devices and Services*, MobileHCI, pp. 55–58, ACM, 2010.

[42] G. Bigwood, F. Ben Abdesslem, and T. Henderson, "Predicting location-sharing privacy preferences in social network applications," in *Proc. of the First Workshop on recent advances in behavior prediction and pro-active pervasive computing (AwareCast)*, June 2012.

[43] C. Dong, H. Jin, and B. Knijnenburg, "Predicting privacy behavior on online social networks," in *Proc. of the Int. AAAI Conf. on Web and Social Media*, 2015.

[44] I. Bilogrevic, K. Huguenin, B. Agir, M. Jadliwala, M. Gazaki, and J.-P. Hubaux, "A machine-learning based approach to privacy-aware information-sharing in mobile social networks," *Pervasive Mob. Comput.*, vol. 25, pp. 125–142, Jan. 2016.

[45] J. Xie, B. P. Knijnenburg, and H. Jin, "Location sharing privacy preference: Analysis and personalized recommendation," in *Proc. of the 19th Int. Conf. on Intelligent User Interfaces*, IUI '14, (NY, USA), pp. 189–198, ACM, 2014.

[46] A. Acquisti and J. Grossklags, "Privacy and rationality in individual decision making," *IEEE Security and Privacy*, vol. 3, pp. 26–33, Jan. 2005.

[47] M. Götz, S. Nath, and J. Gehrke, "Maskit: Privately releasing user context streams for personalized mobile applications," in *Proc. of the ACM Int. Conf. on Management of Data (SIGMOD)*, pp. 289–300, ACM, 2012.

[48] M. Xue, P. Kalnis, and H. Pung, "Location diversity: Enhanced privacy protection in location based services," in *Location and Context Awareness*, vol. 5561 of *Lecture Notes in Computer Science*, pp. 70–87, Springer Berlin Heidelberg, 2009.

[49] B. Lee, J. Oh, H. Yu, and J. Kim, "Protecting location privacy using location semantics," in *Proc. of the 17th ACM SIGKDD Int. Conf. on Knowledge discovery and data mining (KDD)*, pp. 1289–1297, ACM, 2011.

[50] M. L. Damiani, E. Bertino, and C. Silvestri, "The probe framework for the personalized cloaking of private locations," *Trans. Data Privacy*, vol. 3, no. 2, pp. 123–148, 2010.

[51] Z. Riaz, F. Dürr, and K. Rothermel, "Understanding vulnerabilities of location privacy mechanisms against mobility prediction attacks," in *Proc. of the 14th EAI Int. Conf. on Mobile and Ubiquitous Systems: Computing, Networking and Services (MobiQuitous'17)*, (Melbourne, Australia), ACM, 2017.

[52] G. Theodorakopoulos, R. Shokri, C. Troncoso, J.-P. Hubaux, and J.-Y. Le Boudec, "Prolonging the hide-and-seek game: Optimal trajectory privacy for location-based services," in *Proc. of the 13th Workshop on Privacy in the Electronic Society (WPES)*, pp. 73–82, ACM, 2014.

[53] Y. Xiao and L. Xiong, "Protecting locations with differential privacy under temporal correlations," in *Proc. of the 22nd ACM SIGSAC Conf. on Computer and Communications Security*, CCS '15, (NY, USA), pp. 1298–1309, ACM, 2015.

[54] C. Dwork, "Differential privacy," in *33rd Int. Colloquium on Automata, Languages and Programming, part II (ICALP 2006)*, vol. 4052 of *Lecture Notes in Computer Science*, (Venice, Italy), pp. 1–12, Springer Verlag, July 2006.

[55] M. E. Andrés, N. E. Bordenabe, K. Chatzikokolakis, and C. Palamidessi, "Geo-indistinguishability: Differential privacy for location-based systems," in *Proc. of the 2013 ACM SIGSAC Conf. on Computer & Communications Security*, CCS '13, (NY, USA), pp. 901–914, ACM, 2013.

[56] M. Benisch, P. G. Kelley, N. Sadeh, T. Sandholm, J. Tsai, L. F. Cranor, and P. H. Drielsma, "The impact of expressiveness on the effectiveness of privacy mechanisms for location-sharing," in *Proc. of the 5th Symposium on Usable Privacy and Security*, SOUPS '09, (NY, USA), pp. 22:1–22:1, ACM, 2009.

[57] K. Connelly, A. Khalil, and Y. Liu, "Do i do what i say? observed versus stated privacy preferences," in *Int. Conf. on Human-Computer Interaction (INTERACT)*, 09/2007 2007.

[58] Y. Liu, K. P. Gummadi, B. Krishnamurthy, and A. Mislove, "Analyzing facebook privacy settings: User expectations vs. reality," in *Proc. of the 2011 ACM SIGCOMM Conf. on Internet Measurement Conf.*, IMC '11, (NY, USA), pp. 61–70, ACM, 2011.

[59] R. Ravichandran, M. Benisch, P. G. Kelley, and N. M. Sadeh, "Capturing social networking privacy preferences," in *Proc. of the 9th Int. Symposium on Privacy Enhancing Technologies*, PETS '09, (Berlin, Heidelberg), pp. 1–18, Springer-Verlag, 2009.

[60] N. Sadeh, J. Hong, L. Cranor, I. Fette, P. Kelley, M. Prabaker, and J. Rao, "Understanding and capturing people's privacy policies in a mobile social networking application," *Personal Ubiquitous Comput.*, vol. 13, pp. 401–412, Aug. 2009.

[61] L. Fang and K. Lefevre, "Privacy wizards for social networking sites," in *Proc. of the 19th Int. Conf. on World Wide Web*, WWW '10, (NY, USA), pp. 351–360, ACM, 2010.

[62] D. Freni, C. Ruiz Vicente, S. Mascetti, C. Bettini, and C. S. Jensen, "Preserving location and absence privacy in geo-social networks," in *Proc. of the 19th ACM Int. Conf. on Information and Knowledge Management*, CIKM '10, (NY, USA), pp. 309–318, ACM, 2010.

[63] S. Jaiswal and A. Nandi, "Trust no one: A decentralized matching service for privacy in location based services," in *Proc. of the Second ACM SIGCOMM Workshop on Networking, Systems, and Applications on Mobile Handhelds*, MobiHeld '10, (NY, USA), pp. 51–56, ACM, 2010.

[64] S. Pidcock and U. Hengartner, "Zerosquare: A privacy-friendly location hub for geosocial applications," in *Proc. of IEEE Mobile Security Technologies Workshop*, (MoST'13), pp. 1–10, IEEE Press, 2013.

[65] S. Guha, M. Jain, and V. N. Padmanabhan, "Koi: A location-privacy platform for smartphone apps," in *Proc. of the 9th USENIX Conf. on Networked Systems Design and Implementation*, NSDI'12, (Berkeley, CA, USA), pp. 14–14, USENIX Association, 2012.

[66] F. Dürr, P. Skvortsov, and K. Rothermel, "Position sharing for location privacy in non-trusted systems," in *Proc. of the 9th IEEE Int. Conf. on Pervasive Computing and Communications (PerCom 2011)*, pp. 189–196, 2011.

[67] Z. Riaz, F. Dürr, and K. Rothermel, "Optimized location update protocols for secure and efficient position sharing," in *Proc. of the 2nd Int. Conf. on Networked Systems (NetSys 2015)*, pp. 1–8, 2015.