

RESEARCH

Open Access



Where you go is who you are: a study on machine learning based semantic privacy attacks

Nina Wiedemann^{1*}, Krzysztof Janowicz², Martin Raubal¹ and Ourania Kounadi²

*Correspondence:
nwiedemann@ethz.ch

¹ Institute of Cartography and Geoinformation, ETH Zürich, Zurich, Switzerland

² Department of Geography and Regional Research, University of Vienna, Vienna, Austria

Abstract

Concerns about data privacy are omnipresent, given the increasing usage of digital applications and their underlying business model that includes selling user data. Location data is particularly sensitive since they allow us to infer activity patterns and interests of users, e.g., by categorizing visited locations based on nearby points of interest (POI). On top of that, machine learning methods provide new powerful tools to interpret big data. In light of these considerations, we raise the following question: What is the actual risk that realistic, machine learning based privacy attacks can obtain meaningful semantic information from raw location data, subject to inaccuracies in the data? In response, we present a systematic analysis of two attack scenarios, namely location categorization and user profiling. Experiments on the Foursquare dataset and tracking data demonstrate the potential for abuse of high-quality spatial information, leading to a significant privacy loss even with location inaccuracy of up to 200 m. With location obfuscation of more than 1 km, spatial information hardly adds any value, but a high privacy risk solely from temporal information remains. The availability of public context data such as POIs plays a key role in inference based on spatial information. Our findings point out the risks of ever-growing databases of tracking data and spatial context data, which policymakers should consider for privacy regulations, and which could guide individuals in their personal location protection measures.

Keywords: Location privacy, Place labelling, Semantic privacy, Human mobility

Introduction

In the age of big data, an unprecedented amount of information about individuals is publicly available. Not only the information from social media profiles can be exploited to gain rich insights into the private life of individuals, but also data that is collected by applications on-the-fly. Collecting and selling such data has become a business model of commercial consumer data brokers [76], who distribute individual data of users, oftentimes without their awareness [15]. A particularly popular source is location data, as the whereabouts of people allow rich insights into their daily activities [5, 22, 38, 63], for example, for the purpose of profiling. Even though awareness for (location) privacy has increased in recent years [2], this is oftentimes not reflected in user behavior, which

has been termed the “privacy paradox” [7, 68]. Only gradually, companies are reacting to imposed privacy regulations and the efforts of privacy advocates’ groups [29]. For example, AppleTM is giving back control over data sharing decisions in the iPhoneTM, including location data,¹ and StravaTM offers to restrict track-visibility in their app for recording physical activities.²

The simplest way to protect location data is a form of masking or obfuscation of the exact geographic coordinates [47]; i.e., deliberately reducing the data quality [21]. While hiding the exact location may provide some anonymity, the risk of unwanted *semantic* inference from the raw location data remains. Consider the following scenario: a data broker obtains location data from a user (e.g., sold from a smartphone app), with the goal to enrich the data and sell it to other companies. For enrichment, he combines the track points with spatial context data such as public points of interests. For instance, if a user is detected in a busy city district at night, it is very likely that the user is in a bar or club. If the data broker processes location data collected over a longer time period in this fashion, intimate information about the user’s hobbies and interests is unveiled. These user profiles can be sold for targeted advertising, or could even be misused by insurance companies or for influencing elections.³ Following Tu et al. [83], we term this type of unwanted inference “semantic privacy attack”, in contrast to previous work on location privacy that was mainly concerned with user *re-identification* attacks [18, 48, 53, 54, 73].

Here, we aim to quantify the risk of an adversary to *derive meaningful user profiles* from the raw location data of a single user. We argue that a smart attacker would tackle this problem by utilizing spatial and temporal information for categorizing the locations that a user has visited, drawing from methods developed in reverse geocoding [1, 12, 24, 46, 51, 67], activity categorization [17, 25, 62, 66, 75, 85] and place labeling [19, 42, 91, 97] research. For example, if the location data indicates a two-hour stay in a place with many bars nearby, the attacker may derive that the activity falls into the category “Nightlife”. In a second step, the attacker could aggregate the (predicted) categories of all locations that a user visited into a location-based user profile. For example, the profile is 60% “Dining”, 30% “Retail”, and 10% “Nightlife”. In short, we consider the following two semantic attack scenarios:

Task 1: Given a location visit defined by geographic coordinates and a visitation time, the attacker aims to assign the place to the correct category.

Task 2: Given the location visitation pattern of a user, the attacker aims to derive a user profile, defined as the visitation frequencies to each of the location categories.

To the best of our knowledge, this type of location-based user profiling has not been regarded as a privacy attack, and similar definitions for user profiles are mainly found in literature on recommender systems [86, 93]. Note that if these tasks are feasible, the attacker would not only know about activity frequencies but also about when and where each type of activity is preferably carried out. The input data of the attacker is assumed

¹ <https://support.apple.com/guide/iphone/control-the-location-information-you-share-iph3dd5f9be/ios>.

² <https://support.strava.com/hc/en-us/articles/115000173384-Edit-Map-Visibility>.

³ See Appendix for further specific scenarios.

to consist only of geographic coordinates and timestamps. Such data could stem from GNSS tracking data, from Call-Detail-Records [94], or other forms of movement data.

According to Keßler and McKenzie [44], “an individual’s level of geoprivacy cannot be reliably assessed because it is impossible to know what auxiliary information a third party may have access to.” (p. 11). However, one can attempt to quantify the level of privacy by simulating realistic scenarios and measuring the accuracy of the attacker [77, 78]. By realistic, we mean that an attacker tries to enrich the raw data with as much information as possible and employs sophisticated algorithms to analyze patterns in such information. We believe that there is a lack of work analyzing (1) which spatial and temporal information may be exploited, (2) how the data quality, as well as the level of intended inaccuracy due to location protection measures, affects an attacker’s accuracy, and (3) what is the relation to the density and quality of spatial context data, e.g., public POIs. We, therefore, evaluate the effectiveness of machine learning based semantic privacy attacks in different scenarios with respect to the information available to the attacker and, similar to [27], varying the data accuracy by means of random perturbations of the location.

Related work

Reverse geocoding and activity categorization

Many studies utilize a well-known dataset of location check-ins from the Location-based Social Network (LBSN) Foursquare, which is very suitable due to its size, its detailed POI categorization taxonomy, and the availability of user-wise check-in data. The POIs and visitation patterns were analyzed for recommender system applications [92], for deriving interpretable latent representations of venues [3] or to infer urban land-use via clustering of POI data [26]. Yang et al. [89] train models on the Foursquare dataset to infer spatio-temporal activity preferences of users for the purpose of place recommendation. In this work, we take a machine learning viewpoint and regard the Foursquare data as a labeled dataset that is suitable to model the real-life scenario where an attacker aims to categorize the locations of an *unseen* user.

However, it was shown that not only spatial but also temporal information about location visits could be exploited to infer location categories [56]. This has been reported implicitly in other work, for example, Do and Gatica-Perez [19] regard the problem of automatic place labeling into 10 categories, leveraging visitation patterns, e.g., temporal features (start and end time or duration) and visitation frequency from smartphone data. McKenzie et al. [59] and McKenzie and Zhang [57] connect this observation to geoprivacy research by showing that temporal information or texts from social media posts can be exploited for inference about user locations by matching their semantic signatures [41, 58]. While our study is on location categorization and user profiling, in contrast to user localization, their study inspired us to include temporal features in the attack scenario and to contrast their effect on the attacker’s success to the one due to spatial information.

Furthermore, work on user profiling from location data (our second attack task) can mainly be found in the literature on recommender systems, which is surveyed in [6]. The POI embedding of users can be viewed as their location profile, for example, with graph-based embeddings [86]. Ying et al. [93] compare users by their “semantic trajectory”, defined

as the categories of sequentially visited places. We follow their approach but disregard the order of places.

Location privacy research

Privacy risks and potential privacy preservation techniques were studied extensively in the past years [70], including the risks from machine learning [50]. In location privacy research, it was found that a few track points are sufficient to uniquely identify users [18, 30, 73], that it is possible to track people just by the speed and starting location [28] or by accelerometer readings [36], and that even topological representations of movement data without coordinates can be exploited to match users [53, 54]. A common aim of many works is to maintain the performance of a location-based service while providing privacy guarantees; i.e., to optimize the privacy-utility trade-off [9, 80]. Various frameworks for protecting sensitive location data were proposed [10, 21, 43, 60, 61, 74], oftentimes based on k -anonymity [33, 35, 81] or ϵ -differential privacy [4, 14, 20, 23, 37, 40]. For an overview of possible privacy attacks on location data and protection methods we refer to the reviews by Kounadi et al. [47] and Wernke et al. [84].

This work instead analyzes privacy attacks that aim to reveal personal information, i.e., interests and behavioural patterns. Related work in this direction, for example, investigates to what extent demographics (e.g., age or gender) and visited POIs can be derived from location traces [49]. Crandall et al. [16] and Olteanu et al. [64] analyze co-location events and the risk to infer social ties. Tu et al. [83] recently termed the inference of private semantic information from movement trajectories as a “semantic” privacy attack, and they specifically regard contextual POI data as semantics. We build up on their definition and consider attacks that aim to infer POI categories. Tu et al. [83] propose l -diversity and t -closeness measures to protect trajectories from semantic inference. However, these approaches rely on trusted third-party (TTP) services that mask the data of multiple users and update their data iteratively in online applications [45, 65]. Omitting the dependence on a TTP is possible, for example, with simple location obfuscation methods, i.e., adding random noise to coordinates or methodologically translating geographic coordinates in space [4, 21]. Zhang et al. [96] and Götz et al. [31] further propose context-aware masking techniques that are applicable to new users, and Qiu et al. [69] propose a framework for obfuscating trajectory semantics. Here, we do not aim to compare location protection methods, but to quantify the risks of realistic semantic privacy attacks without access to a TTP service. Thus, we utilize location obfuscation mainly as a tool for modelling reduced data quality in real-world scenarios. As proposed by Shokri [77], we evaluate the attacker’s accuracy to quantify privacy loss.

Experimental design

We take a machine learning viewpoint and assume that the attacker aims to learn a mapping from visited locations to categories. The available data are a time series of location visits of a new user u . We group the raw data by location in order to gather temporal information about the visitation patterns to one location. The dataset D_u for one user u can be formalized as

$$D_u = \{(l_i^u, [t_1(l_i^u), t_2(l_i^u), \dots]) \mid l_i^u \in L_u\} = \{(l_i^u, T_u(l_i^u)) \mid l_i^u \in L_u\} \quad (1)$$

where L_u is the set of all locations visited by the user u , l_i^u is one location in L_u , and $t_j(l_i^u)$ is the time of the j -th visit of user u to location l_i^u . For simplicity, we abbreviate the ordered list of visit times as $T_u(l_i^u)$. Furthermore, we assume there exists an unambiguous mapping $c : L \rightarrow C$ from each location to a category from a predefined location-category set C . For example, $C = \{\text{Dining, Sports, Shopping}\}$ and the categories for user u are $c(l_1^u) = \text{Shopping}$, $c(l_2^u) = \text{Dining}$, etc.

The attacker aims to learn a model \hat{c} that approximates the true mapping c . The most straightforward approach for \hat{c} is a spatial nearest neighbor join with a public POI dataset; i.e., if the spatially closest POI is a restaurant, then $\hat{c}(l_i^u) = \text{Dining}$. More sophisticated methods could pool the spatial and temporal information and frame \hat{c} as a machine learning model. Here, we simulate the latter via the XGBoost (XGB) algorithm [13]. XGB is a tree-based boosting method that was repeatedly shown to outperform Neural Networks on tabular data [32] and is known to perform particularly well in classification tasks with unbalanced data, as it is the case here. We also chose XGB for its interpretability and since it was empirically superior to a multi-layer perceptron approach in our tests (see section “[Machine learning model](#)”).

Together, we consider the following attack scenarios:

- Spatial join: For each user-location l_i^u , the category of the public POI that is closest to its geographic location $(x(l_i^u), y(l_i^u))$ is assigned.
- XGB temporal: The attacker employs a learning approach, namely XGBoost, based on temporal information derived from $T_u(l_i^u)$ (see section “[Temporal features](#)”).
- XGB spatial: The attacker trains a model on spatial context features (see section “[Spatial features](#)”). No temporal visit information is considered, only coordinates and publicly available POI data.
- XGB spatiotemporal: The model is trained on all available features, i.e., features derived from $(x(l_i^u), y(l_i^u))$ and $T_u(l_i^u)$ as well as available POI data.

In addition, we report the results for an uninformed attacker, where the predictions are drawn randomly from a categorical distribution, with the class probabilities corresponding to the class frequency in the training data.

In our experimental setup, we take an ML perspective and simulate the attack on *new* users via a train-test data split. Evaluating the accuracy of this attack requires a *labeled* dataset \mathcal{D} of user-location pairs l_i^u ; i.e., the location category $c(l_i^u)$ must be *known*. GNSS tracking datasets usually do not provide detailed and reliably place labels. Instead, we found a public dataset from the location-based social network Foursquare most suitable for this experiments since location visits are given as check-ins to places of known categories. The dataset was already used for related tasks [3, 26, 89, 92], but without regarding privacy aspects. The places are categorized into 12 distinct classes according to the Foursquare place taxonomy (see Fig. 3 for the list of categories and section “[Data and preprocessing](#)” for details). Additionally, we also use the Foursquare places as public POI data that may be exploited by the attacker as auxiliary spatial context data. Figure 1 provides a visual overview of the experimental setup. The input data (geographic

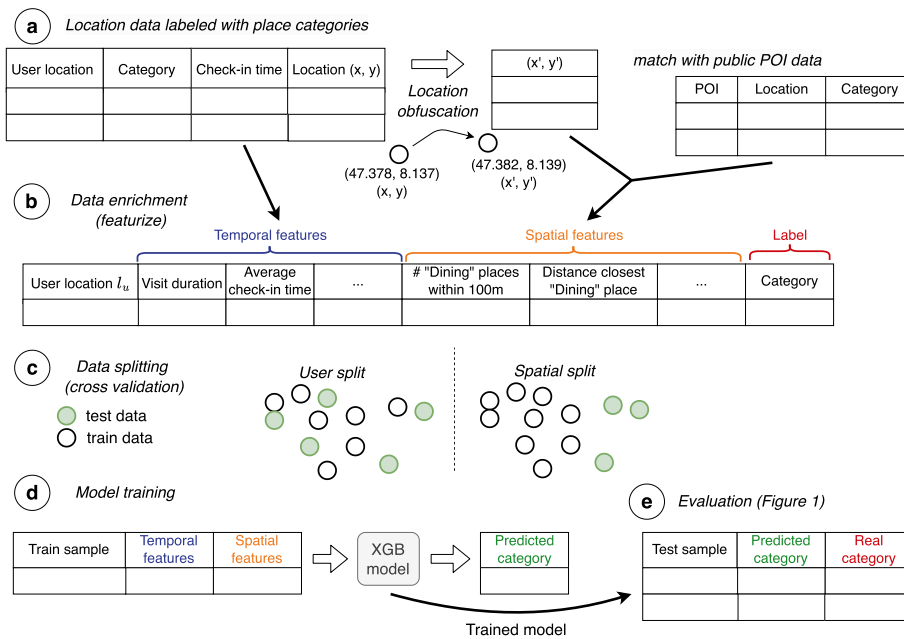


Fig. 1 Overview of the experimental setup. The samples are spatiotemporal data about location visitation patterns. We simulate reduced data quality and potential protection measures by obfuscating the geographic coordinates (a). The samples are then featurized into vectors encoding temporal visitation patterns and spatial context (b). We simulate a privacy attack on new users by a train-test split (c) and train an XGB model to predict the location category (d). The accuracy is evaluated on the test data (e)

coordinates and time points) are enriched with spatial and temporal features. Before computing spatial features, the location is *obfuscated* within a varying radius r to simulate GNSS inaccuracies and possible privacy protection measures (see section “[Location masking](#)”). Then, the data is split into train and test sets, either by user or spatially, to simulate transfer to new users or even to other geographic regions. All results are reported on the combination of all test sets from tenfold cross validation (see section “[Data split](#)”).

Results

Effect of location obfuscation on place labeling accuracy

The results for task 1 (location categorization) are evaluated in terms of accuracy, i.e., the number of correctly categorized places divided by the total number of samples, across all users and all locations (90,790 samples in NYC and 211,834 in Tokyo):

$$Acc(\hat{c}, c) = \frac{\sum_{I_i^u \in \mathcal{D}} \mathbb{1}[\hat{c}(I_i^u) = c(I_i^u)]}{|\mathcal{D}|} \tag{2}$$

Figure 2 shows the classification accuracy of the attack scenarios by the obfuscation radius. Note that $r = 0$ is an unrealistic scenario, since the check-in data and the public POI context data are both from the Foursquare dataset and are based on the exact same set of geographic coordinates. Thus, a simple spatial nearest neighbor join of the check-in location with public POIs achieves 100% accuracy if no obfuscation is applied. Deriving a user’s location from tracking data would obviously hardly yield the exact same point coordinates as a public POI. We, therefore, consider more realistic scenarios with




Fig. 2 Effect of location obfuscation radius on the attacker’s performance in categorizing locations. Spatial information are valuable for an ML algorithm even with up to 1 km of obfuscation

weak obfuscation, and, additionally, protective scenarios with strongly obfuscated coordinates. Figure 2 shows that the accuracy decreases rapidly with the obfuscation radius, but even when the attacker uses only temporal information, the accuracy is 39.1% for Tokyo and 29.7% for NYC, which is significantly better than random (grey line). On top of that, spatial context information can benefit the attack even when the location is obfuscated within a radius of 1 km. This is remarkable and demonstrates the danger of powerful privacy attacks that make use of public POI data. In the appendix, we relate these findings to the spatial autocorrelation of place types (Fig. 17) and we demonstrate that the results of NYC and Tokyo are surprisingly similar (see Appendix Fig. 12). Furthermore, the categorization accuracy depends on the place type; i.e., some categories are harder to detect than others. Figure 3 presents the confusion matrix for the attack scenario at 100 m obfuscation. The error is more evenly distributed over categories than expected, although “Dining” and “Retail” are predicted disproportionately often (see Appendix Fig. 11).

Figure 2 additionally compares a user split to a spatial split to analyze generalization across space (see section “Data split”). Note that a user split is expected to be strictly better than the spatial split because the input data does not include user-identifying information such as age or gender, rendering the generalization to new users as easy as to any new samples. Surprisingly, the spatial cross-validation split only has a minor effect on the attacker’s accuracy (decrease of $\sim 5\%$). We conclude that the attacker’s training data set is not required to cover the exact same region for the privacy attack to be successful.




Fig. 3 Normalized confusion matrix of predictions in NYC with Foursquare data and location prediction with an obfuscation radius of 100 m. The accuracy is rather balanced across categories; however, many activities are erroneously classified as “Dining”

User profiling error for probabilistic and frequency-based profiling

While the ability of a potential attacker to categorize visited locations is concerning, we argue that the main risk is *user profiling* based on the predicted categories. It is unclear to what extent the high categorization accuracy on a location level transfers to a high profiling accuracy on a user level. Here, we define a user profile as the frequency of different types of locations in the user’s mobility patterns. Our definition corresponds to the term-frequency in the TF-IDF statistic,⁴ which measures the frequency of a word in a specific document in relation to the overall occurrence of the term (in the corpus). Here, the “words” are place categories and a “document” is the location trace of one user. We provide examples for such TF-based user profiles in Fig. 4b (“Ground truth”). In the following, we define $p(u)$ as the profile of user u , and $p_c(u)$ as the entry of the vector corresponding to the frequency of category $c \in C$. For example, the ground truth profile of User 1 in Fig. 4b corresponds to $[0.25, 0.5, 0.25]$, since $p_{\text{Dining}}(\text{User 1}) = 0.25, p_{\text{Retail}}(\text{User 1}) = 0.5, p_{\text{Nightlife}}(\text{User 1}) = 0.25$. In this study, we aim to quantify how accurately the adversary could predict $p(u)$. The evaluation of user profiling performance boils down to comparing the difference between two categorical distributions, namely the distributions of the real profile $p(u)$ versus the predicted category frequencies $\hat{p}(u)$:

$$E_{\hat{p}(u), p(u)} = \sqrt{\sum_{c \in C} (\hat{p}_c(u) - p_c(u))^2} \tag{3}$$

⁴ The inverse term frequency (IDF) would correspond to a weighting of the user’s category-frequency by the overall frequency of this category in the data, giving higher weights to rare categories. Since the weights are the same for all users, IDF does not help to distinguish users, neither intuitively nor empirically. We therefore only characterize users by the easily interpretable TF term.

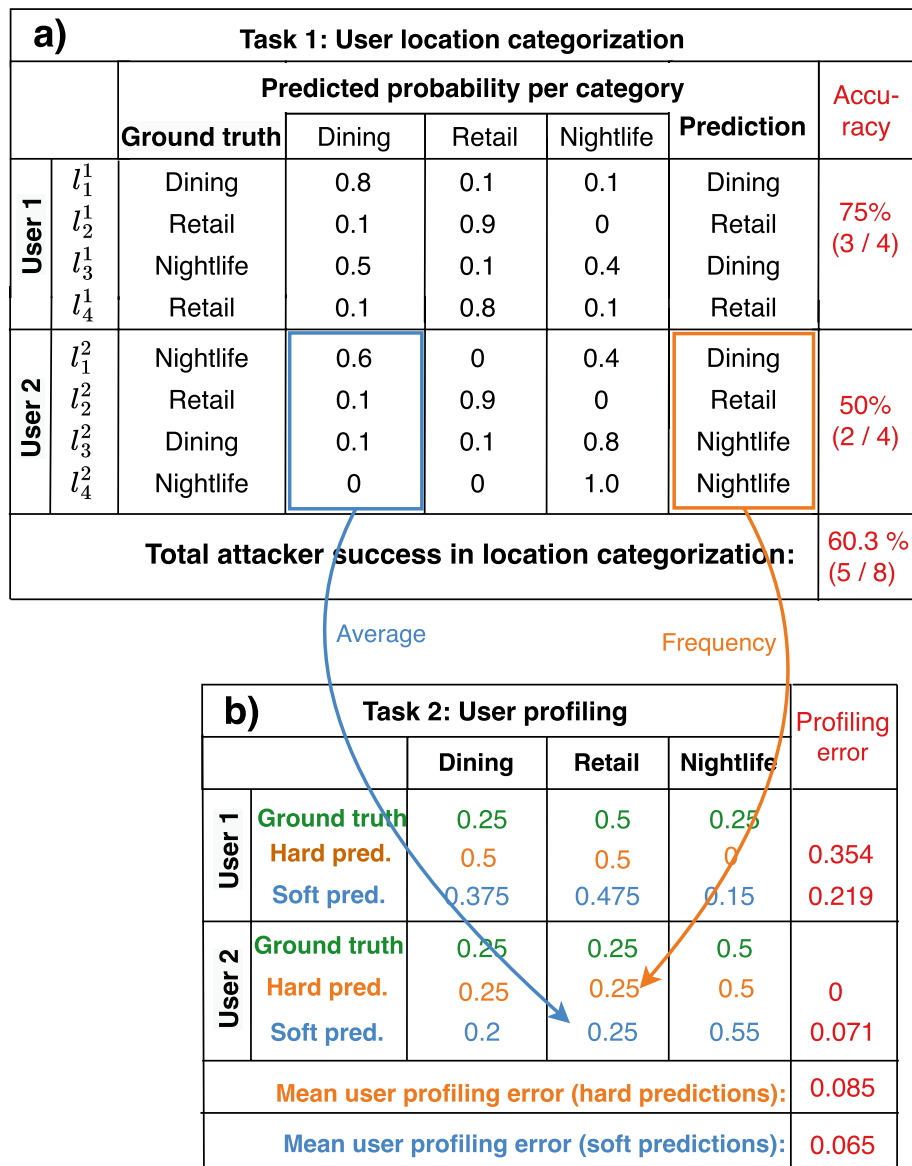


Fig. 4 User profiling from location labelling. **a** The true location category is compared to the category with the highest predicted probability. The place labelling accuracy is computed as the ratio of categories where the prediction matches the ground truth. **b** The predicted labels for individual location visits can be aggregated per user to yield an estimated user profile; reflecting behaviour and interests. The visits are aggregated either by their frequency per category (orange) or by their average predicted probability (blue). The profiling error expresses the difference between the predicted profile and the true profile

The attacker can estimate the profile $\hat{p}(u)$ simply by counting the predicted place categories. For example, in Fig. 4 the “Retail” category is predicted one out of four times for user 2 and therefore takes a value of 0.25 in the profile (see orange arrow). However, many ML-based classification models actually predict a “probability”⁵ for each category,

⁵ The probability distribution over categories is usually derived from the predicted values with a softmax function or by averaging hard predictions of base estimators and is, therefore, by no means the actual posterior distribution. While the provided uncertainties are oftentimes poorly calibrated [34], they nevertheless add information to the final predicted label.




Fig. 5 Comparison of user-profiling errors achieved from averaging “hard” predictions or “soft” prediction probabilities for each category. Probabilistic classifications improve the spatial attack, in particular for lower-quality location data

as shown in Fig. 4a. The XGBoost model, for example, outputs the prediction frequency of each category among its base learners (decision trees). Probabilistic predictions provoke a second way to estimate $\hat{p}(u)$, namely by averaging the predicted probabilities per category (see blue arrow in Fig. 4). In the following, we term the first option (computing the frequency of predicted categories, orange) as “hard” profiling and the second option (averaging category-wise probabilities, blue) as “soft” user profiling. As shown in the toy example in Fig. 4, soft profiling can increase or decrease the error compared to hard profiling (e.g., decrease from 0.354 to 0.219 for user 1, but increase from 0 to 0.071 for user 2).

In Fig. 5, we empirically compare both strategies on our dataset in terms of the error E defined above. Only the error for the strongest attack scenario (XGB spatio-temporal) is shown, averaged over cities (NYC and Tokyo). The profiling error is significantly lower for the soft profiling strategy that is based on probabilistic predictions. In particular, the error of “hard” profiling increases proportionally with a doubling of the obfuscation radius, while the error of soft-labeling increases sub-linearly (see Fig. 5). This result is consistent for all considered scenarios. It demonstrates that well-calibrated probabilistic prediction methods are more dangerous in terms of user profiling than point predictors, even if the latter may achieve a higher place classification accuracy.

All further results are reported for the *soft* predictions in order to simulate the strongest attack.

User reidentification accuracy based on the estimated profiles

Judging from the error alone it is difficult to interpret how much the user profile actually reveals. Such interpretation depends on the variance of the user profiles: For example, if all users have the same profile, the prediction error may be very low, but there is no value in profiling. As a more interpretable metric, we follow previous privacy research




Fig. 6 User-profiling performance of different semantic attacks, in terms of the top-5 accuracy of re-identifying users by their profile. With an obfuscation radius of around 400 m, the user profiling accuracy converges to zero

and analyze the possibility of re-identifying users by their predicted profile. Given the pool of ground-truth user profiles (Fig. 4b green), we match the predicted profiles by finding their nearest neighbors in the pool based on the Euclidean distance of their profile vectors. We report the results in terms of top-5 re-identification accuracy, also called hit@5.

In Fig. 6, the re-identification accuracy is shown by the attack scenario. A corresponding plot of the profiling error is given in the appendix (Fig. 14). Although the accuracy decreases quickly with stronger obfuscation, it is still larger than 10% even with an obfuscation radius of 1.2 km. The average uninformed (random) identification accuracy is 0.6% on average, with 1083 users in NYC and 2293 users in Tokyo. To compare the decay of the user profiling performance to the decay in place categorization accuracy (Fig. 2), we fit an exponential function of the form $f(x) = a + c \cdot e^{-x \cdot \lambda}$ to both results. The place categorization accuracy decays with $a = 0.3439, \beta = 0.0097, c = 0.6216$, indicating that the accuracy decreases with a rate of $e^{-0.0097} = 0.9903$ but converges to around 0.3439. The function fit for the user identification accuracy yields $a = 0.0625, \beta = 0.0121, c = 0.9518$. In other words, with every 50 m added to the location obfuscation radius, the user re-identification accuracy is reduced by a factor of 0.5488 ($= e^{-0.0121 \cdot 50}$). At an obfuscation radius of $r = 57.43$, the accuracy has approximately halved. This firstly demonstrates that place categorization does not directly translate into user profiling, as the profiling accuracy decays faster than the categorization accuracy, and secondly gives guidance for selecting a suitable masking radius.

Induced privacy loss of ML-based privacy attacks

Finally, we transform the re-identification accuracy into a *privacy loss* metric following [53]. They define the privacy loss PL for one user $u \in U$ as

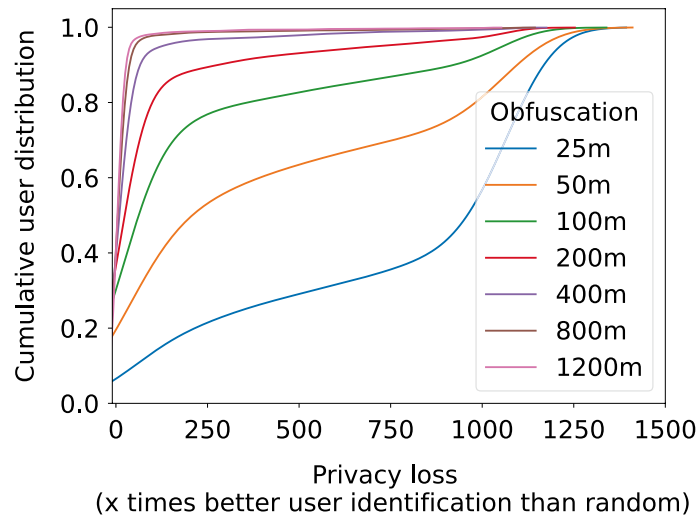


Fig. 7 Cumulative distribution of the privacy loss per user caused by the strongest attack scenario

$$PL(u) = \frac{P_{attack}(u = u^* | D_u)}{P_{uninformed}(u = u^*)} \quad (4)$$

where $P_{uninformed}$ is the probability of an uninformed adversary to match u to the true user u^* , corresponding to a random pick from all users U , so $P_{uninformed} = \frac{1}{|U|}$. The probability of an informed adversary, on the other hand, is the probability to match the user to the correct profile by utilizing sensitive user data including geographic coordinates and visitation times. We assume that given a pool of users U , the attacker would match u to a user $u_i \in U$ from the pool with a probability proportional to the similarity of their profiles:

$$P_{attack}(u = u_i | \mathcal{D}) \propto \text{softmax}(sim(u, u_i)) = \frac{e^{sim(u, u_i)}}{\sum_{j=1}^{|U|} e^{sim(u, u_j)}} \quad (5)$$

where we define the similarity as the inverse distance of the user profile vectors $sim(u, u_i) = (E_{\hat{p}(u), \hat{p}(u_i)})^{-1}$. Note that Manousakas et al. [53] use a rank-based measure of similarity, which however seems unintuitive given that we know the exact distance between each pair of user-profiles and not only their respective rank.

The median privacy loss is 11 if the adversary is given spatio-temporal information where the locations are obfuscated by 100 m (see Appendix Table 1). In other words, the adversary is still 11 times better at re-identifying a user by his profile than with a random strategy. Moreover, the adversary with spatio-temporal data is 9.9 times better than an adversary that uses only temporal information, even though the spatial data are obfuscated up to 100 m. At higher location obfuscation, the privacy loss converges. The strongest attack only yields a median privacy loss of 3.74 at 200 ms obfuscation radius and 2.13 at 400 m. However, the privacy loss strongly varies across users. Figure 7 shows the cumulative distribution of users. If the locations are obfuscated by 100 m, around 80% of the users have a privacy loss lower than 250; however, the distribution is heavy-tailed with a considerable number of users that are still easy to identify. Nevertheless,




Fig. 8 Feature importances in the XGBoost classifier. The occurrence of different categories and their mean distance are the most important features for place categorization

we conclude that obfuscating the location with a radius between 100 and 200 ms would significantly reduce the risk of successful profiling attacks for a large majority of users.

Features that affect the predictability of place categories

One advantage of boosted-tree based machine learning methods such as XGBoost is that decision trees are interpretable. While the individual decision boundaries are not transparent in large ensembles of trees, one can still compute the importance of individual features in terms of their mean decrease of data impurity. The respective importance of the spatial and temporal features included in our study are shown in Fig. 8. The most important spatial features are the number of POIs per category among the *k* nearest POIs. The spatial embedding features derived with the space2vec (embed 0–embed 16) method apparently do not add much information. The time of the day, expressed in sinus and cosinus of the hour and binary variables for morning, afternoon and evening, also play a significant role, highlighting the relevance of temporal information.

Dependency on POI data quality

To simulate incomplete POI data, we subsample 75% or 50% randomly from the Four-square POIs. Furthermore, the performance with POI data from OSM instead of Four-square is evaluated. In this experiment, only the predictions of the strongest attack (XGB spatio-temporal) on NYC check-in data are evaluated. Figure 9 depicts the results, where “Foursquare (all)” corresponds to the results in Fig. 2. The removal of Foursquare POIs has surprisingly little effect on the user identification accuracy. Even with 50% of the POIs, 84.8% of the check-ins can be classified correctly (see Appendix Fig. 13), translating to a top-5 identification accuracy of 94%. This is due to the spatial autocorrelation between places of certain categories (see Appendix Fig. 17).

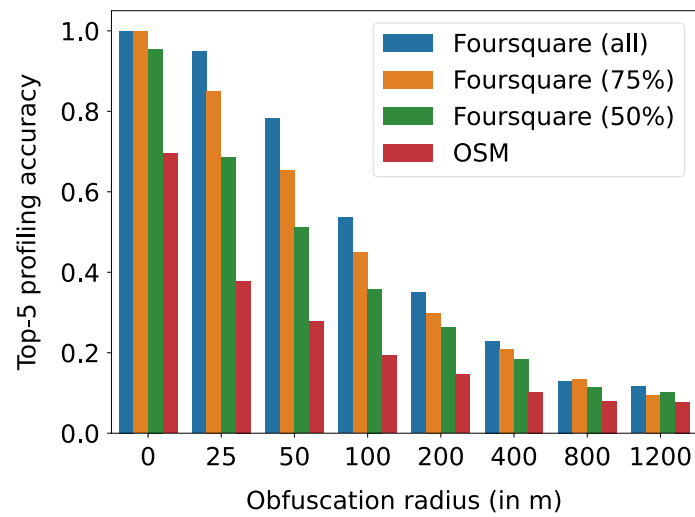


Fig. 9 Dependency of the attacker’s success on the POI quality. The strongest attack scenario based on spatio-temporal data is shown. While the completeness of POI data has a disproportionally low impact on user profiling, using OSM data decreases the attacker’s success




Fig. 10 Place categorization accuracy by POI density (number of POIs within 500 m). Visited places in very dense areas are harder to classify

Meanwhile, it is much harder to classify the category of Foursquare check-ins with OSM POIs. We hypothesize that this is due to substantial differences between OSM and Foursquare POI data. Previous work [95] tried to match cafes in the OSM dataset to cafes in the Foursquare set and find that only around 35% can be matched exactly (Levenshtein distance of labels=1), with a spatial accuracy of around 30–40 m. In addition to these location differences, in our case there are also differences in the place categories, which we partly had to assign *manually* to the OSM POIs (see “Data and preprocessing”). Nevertheless, the low performance with OSM data unveils important difficulties for an attacker to utilize inaccurate, incomplete and dissenting datasets of POIs.

Influence of the POI density

Furthermore, the difficulty level of the attack depends on the density of spatial context data, since it is easier to match a location to a nearby POI if the number of nearby POIs is low. We quantify this relation by computing the number of surrounding POIs within 200 m for all considered places in NYC and Tokyo. In Fig. 10, the place labelling accuracy is shown by POI density groups. Places in dense areas; i.e., with many surrounding POIs, are harder to classify. For example, when the obfuscation radius is 100 m, the mean number of POIs within 200 m around the (non-obfuscated) location is 58 for correctly predicted samples, but 85 for erroneously classified samples. However, the variance between the curves shown in Fig. 10 is lower than expected. Only points with less than ten nearby POIs are significantly easier to match.

The dependence of the predictability on the POI density calls for a context-aware protection scheme [4, 96]. We implement such scheme by setting the obfuscation radius r for a specific location such that at least m public POIs lie within the radius. For the sake of comparability, we tune m to a value that leads to an average obfuscation radius of 200 m ($m = 16$). In other words, when obfuscating each location l within a context-aware radius $r(l)$ that covers exactly 16 public POIs, then $\frac{1}{|\mathcal{D}|} \sum_{l \in \mathcal{D}} r(l) \approx 200$. As desired, this masking scheme destroys the relation between POI density and accuracy. However, our experiments show that the average accuracy *increases* compared to the accuracy reported for location-independent masking in Fig. 2 (accuracy of 0.52 compared to 0.49 for the experiment on NYC-Foursquare data with XGB spatio-temporal). This also holds at a user-level, where the user-profiling performance is higher with context-aware location obfuscation (0.27 vs. 0.23). It seems that the weak obfuscation of locations in high-density regions has a greater effect than the strong obfuscation of isolated places. We conclude that simple context-aware obfuscation based on POI density is not sufficient to reduce privacy risks, at least not at the same average obfuscation level. While the evaluation of protection methods is out of the scope of this work, further work is needed to understand their effectiveness against undesired user-profiling.

Discussion

We have quantified the risks of undesired user profiling in different attack scenarios, varying (1) the information available to the attacker, (2) the location data quality in terms of obfuscation radius, and (3) the POI data quality. We comment on each aspect in the following.

First, our experiments reveal that machine learning methods can efficiently exploit spatial context data, even with low data quality or incomplete data. We further confirm previous findings by McKenzie and Janowicz [56] that even only temporal information about location visits poses a significant privacy risk. This risk may be further increased, for example, if also the opening times of surrounding POIs are used as input features [88]. In general, more powerful ML methods may increase privacy risks beyond our results. A particularly interesting finding is the superiority of *probabilistic* predictions for deriving user profiles. In other words, a potential attacker can estimate the importance of different place types in a user's life without knowing the category for each individual place exactly.

Furthermore, we took a user-centric viewpoint and derived location protection recommendations. The exponential decay of user identification accuracy demonstrates the high effectiveness of simple protective measures, and the results suggest that the privacy risks become negligible when the location is obfuscated with a radius of around 200 m. While such inaccuracy may be intolerable in navigation apps, it yields a good trade-off in other applications such as social media, where the approximate location is still interesting to friends but not yet informative for profiling attacks. It is worth noting that our findings only testify an exponential decay in the quality of user activity profiles, whereas other privacy risks such as user re-identification based on a set of visited points or areas [18] may remain.

However, further experiments on other datasets are necessary to validate the results. Our analysis is based on an experimental setting where each visited location can for sure be matched to a public POI. An attack that aims to classify user activities that are not related to public POIs is, therefore, expected to be more difficult (e.g., detecting a visit to a friend's place). In the appendix (Fig. 15), we provide a study on a GNSS-based tracking dataset where stay points are labeled with a few broad activity categories, but it would be highly interesting to reproduce our results on a GNSS dataset with more detailed place categories. However, datasets that are large and labeled at the same time are rare [11]. Finally, we see a strong dependency of the attacker's success on the density and completeness of spatial context data. Thus, future privacy protection algorithms should not only regard past studies on protection efficiency, but also improvements in public databases. We hope to inspire future research on the risks and, importantly, on suitable protection methods against such novel semantic privacy attacks. Further analysis may, for example, investigate which users are particularly easy or hard to profile. The classification of users into a predefined set of profiles or a cluster of profiles could provide further insights into the actual dangers of unwanted behavior analysis. Finally, it may be an interesting endeavor to develop location protection techniques that specifically target the weaknesses of machine learning models, similar to adversarial attacks [39].

Conclusion

Semantic privacy deserves more attention in geoprivacy research, considering the business case of data brokers and the interest of companies in semantic information in contrast to raw data. Our analysis is a first step towards a better understanding of the actual risk for a user to reveal sensitive behavioral data when sharing location data with applications. Spatial and temporal patterns in location data lead to a significant opportunity for user profiling, even if the coordinates are not accurate. However, this effect diminishes with stronger location protection. Our analysis, therefore, enables users and policy-makers to derive recommendations on a suitable protection strength.

Methods

In the following, our methods are described in detail. Our implementation is available open-source at https://github.com/mie-lab/trip_purpose_privacy.

Data and preprocessing

Check-in data from Foursquare

Our study mainly uses data from the location-based social network *Foursquare*. In contrast to tracking datasets or data from other social networks (e.g., tweets), the Foursquare dataset offers labeled and geo-located place visitation data. Specifically, users check-in at venues, e.g., a restaurant, and the geographic location of the venue as well as a detailed semantic label, e.g., “Mexican restaurant”, are known. Similar to other studies [89, 90], we use the Foursquare subset of New York City and Tokyo in order to simplify location processing and to study the variability of the results over two different cities. The data was collected by Yang et al. [89] from 12 April 2012 to 16 February 2013 and was downloaded from their website.⁶ Note that Foursquare has changed over the years, and the data thus differs from today’s usage of this LBSN. This is not an issue for our study, as the underlying location visitation patterns are expected to remain similar.

As a first step, we clean the category labels of place check-ins of users. We focus on leisure activities and do not consider home and work check-ins for several reasons: (1) Home and work location can be inferred by *temporal* features such as the time of the day and visit duration. Spatial POI data are not necessary. (2) Identifying home and work is possible with simple heuristics, e.g., assigning the most often visited location as home and the second-most-frequently visited location as work. We believe that previous attempts on this task mainly suffer from insufficient data quality and the lack of reference data, and not the difficulty of the task itself. (3) Many Foursquare users in the dataset do not check-in at home or work since the social network was mainly used to share leisure activities, at least in 2012 when the data was gathered and before changes were made to their (check-in) app.

In total, the Foursquare POIs in NYC and Tokyo are labeled with 1146 distinct categories. A taxonomy is provided with 11 groups on the highest level, such as *Dining and Drinking* or *Arts and Entertainment*. We use this categorization as the ground-truth location categories, but make a few changes in order to sufficiently distinguish common types of leisure activities that are relevant for user profiling. Specifically, we divide the category *Dining and Drinking* into categories *Dining* (all kinds of restaurants), *Nightlife* (bars), and *Coffee and Dessert*, based on the label given on lower levels of the taxonomy. Furthermore, the category *Community and Government* is split into the categories *Education* and *Spiritual Centers*. Other subcategories that can not be fitted into these two, e.g. “government building” or “veteran club”, are omitted. Finally, there are around 100 labels in the NYC-Tokyo Foursquare dataset from 2012 that do not appear in the (up-to-date) Foursquare POI taxonomy. We manually assign these labels to categories. The final distribution of the labels in NYC check-ins is shown in Appendix Fig. 16a. For comparison, we additionally experiment with a coarser category set with only six place types. Figure 18 in the appendix demonstrates that the place labelling accuracy increases due to this simplification, but at the cost of less informative user profiles.

Furthermore, the check-in dataset is cleaned by merging subsequent check-ins of the same user at the same location. A check-in event is deleted if it occurs within 1 h of the

⁶ <https://sites.google.com/site/yangdingqi/home/foursquare-dataset?pli=1>.

previous check-in at that location, leading to the removal of 0.496% of the NYC check-ins and 0.63% of the ones in Tokyo.

Public POI data

We assume that the attacker can access public POI data, such as the POIs from Foursquare. However, categorizing check-in locations in the Foursquare data is easy when the Foursquare POIs are given since they correspond exactly in their geographic location and each check-in can (in theory) be matched to a known POI. Apart from obfuscating the check-in location to simulate inaccurate GNSS data, we also simulate incomplete POI data by sampling 50% and 75% of the Foursquare POIs at random.

Last, we simulate a situation with substantially different POI data by using POIs from OSM. The Python package `pyrosm` [82] is used to download all places of the categories “healthcare”, “shop”, “amenity”, “museum”, “religious”, “transportation”, and “station” (public transport) from OSM. The “amenity” category in particular contains a large collection of places, and we first delete all places labeled as “parking space” since they accounted for a large fraction of the data and are irrelevant to our analysis. We further manually re-label the POIs in order to assign place categories. The same categories as in the Foursquare dataset are used and the mapping from OSM-POI-types to our categories is given in detail in our code base.⁷

Spatial and temporal input features to machine learning model

Temporal features

Temporal features are computed from $T_u(l_i^u)$ as the following:

- Visit frequency features: The absolute visit frequencies of location l_i^u , corresponding to $|T_u(l_i^u)|$, and the relative frequency with respect to all check-ins by u , formally

$$f_{\text{visit_frequency}}(l_i^u) = \frac{|T_u(l_i^u)|}{\sum_{l_i^u \in L_u} \sum_{t_j \in T_u(l_i^u)} t_j} \quad (6)$$

The absolute frequencies are scaled with a logarithm to reflect well-known power-law properties of location visitation patterns [8, 72].

- Duration features: In the Foursquare dataset used as training data, the check-outs of location visits are not provided, so only the start time is known. Thus, we approximate the visit duration by computing the time until the next check-in. Since no check-outs are (publicly) available, there are many outliers with gaps over more than a day. We flatten these outliers by scaling logarithmically, and finally, we take the average over the individual visit durations. Formally, the visit time is subtracted from the time of its subsequent check-in, given as the minimum time of all following check-ins of the user:

⁷ https://github.com/mie-lab/trip_purpose_privacy/blob/main/data/osm_poi_mapping.json.

$$f_{\text{dur}}(l_i^u) = \frac{1}{|T_u(l_i^u)|} \sum_{j=0}^{|T_u(l_i^u)|} \log \left(\begin{array}{c} \min_{k,m} t_m(l_k^u) - t_j(l_i^u) \\ \text{s.t. } t_m(l_k^u) > t_j(l_i^u) \end{array} \right) \quad (7)$$

The duration of the last check-in overall is omitted. Although this approximation is very rough due to the dependence on the LBSN usage frequency of users, we empirically observed that it is still helpful for inference.

- Daytime features: Last, the start time is represented by a variety of features: Binary features to indicate whether it is on the weekend, in the morning (before 12 pm), in the afternoon (12 pm–5 pm), in the evening (5 pm–10 pm), or at night (10 pm–midnight). The time thresholds were selected to reflect different activities (e.g. dining vs nightlife). The exact daytime was encoded with trigonometric functions (sine and cosine) to reflect their cyclical properties, as is common in machine learning.

Spatial features

The attacker can utilize the recorded geographic coordinates to predict the location category. However, inputting the raw coordinates to a model is not advisable as they suffer from uncertainty and, more importantly, the model would not generalize to other spatial regions. Thus, spatial features are usually derived from the context of the spatial location, here public POI data, since the categories of surrounding POIs are a valuable predictor [87] for the user’s location category. POI data are, for example, available from the public Foursquare API or from Open Street Map (OSM). In either case, the dataset includes geographic point data and a categorization taxonomy of broad and more specific POI labels, e.g., a POI may be part of both the “Shoe Store” and the overarching “Retail” category. For most spatial features, we only use the broadest level and denote its categories as $\Psi = \{\psi_1, \dots, \psi_n\}$. A POI p has a set of coordinates $(x(p), y(p))$, and is assigned to a main POI category, $c_p(p)$.⁸ For example, p may be assigned to $c_p(p) = \psi_2 = \text{Retail}$.

In the literature, different approaches have been used to extract features from the POI distribution around a specific point. We found empirically that a combination of the following methods yields the best results for our attacker’s task:

- Category-count of the k-nearest POIs: Given a location $(x(l_i^u), y(l_i^u))$, the k closest POIs p_1, \dots, p_k are found via a ball tree search, and the count of each category among those is computed. The result is a feature vector where the first element corresponds to the number of occurrences of the *first* category among the k closest POIs and accordingly for the other categories; formally

⁸ Note that our notation explicitly distinguishes *location* categories ($c(l_i^u) \in C$) from *POI* categories ($c_p(p) \in \Psi$), since they may be different. For example, an attacker could use POI data with 10 categories ($|\Psi| = 10$) to classify user location data into only three categories such as $C = \{\text{Work, Leisure, Eating}\}$.

$$\left[\sum_{i=1}^k \mathbb{1}[c_p(p_i) = \psi_1], \sum_{i=1}^k \mathbb{1}[c_p(p_i) = \psi_2], \dots \right] \quad (8)$$

Furthermore, as an indicator of the POI density at (x, y) , the mean distance from the k nearest POIs is extracted as a feature. We set $k = 20$ in our experiments.

- Count and distance of POIs within a fixed radius: The semantic attack requires more specific distance information of the POIs for each category. For example, if there is no restaurant within 1 km, it is unlikely that the location category is “Dining”. Thus, we consider all POIs around $(x(l_i^u), y(l_i^u))$ within a specified radius r , denoted as the set $P(x, y, r)$,⁹ and again compute the count of each category.

$$\left[\sum_{p \in P(x, y, r)} \mathbb{1}[c_p(p) = \psi_1], \sum_{p \in P(x, y, r)} \mathbb{1}[c_p(p) = \psi_2], \dots \right] \quad (9)$$

In addition, we consider the minimum distance of POIs of one category to the location:

$$\left[\min_{\substack{p \in P(x, y, r) \\ c_p(p) = \psi_1}} \left\| \begin{pmatrix} x \\ y \end{pmatrix} - \begin{pmatrix} x(p) \\ y(p) \end{pmatrix} \right\|, \min_{\substack{p \in P(x, y, r) \\ c_p(p) = \psi_2}} \left\| \begin{pmatrix} x \\ y \end{pmatrix} - \begin{pmatrix} x(p) \\ y(p) \end{pmatrix} \right\|, \dots \right] \quad (10)$$

We set the radius to 200 m based on the results of preliminary experiments. If a category does not appear within the radius, we fill the corresponding vector field by the radius r . As an example, consider that three POIs are found within radius $r = 200$ m of the location: p_1 of category ψ_3 with 50 m distance, p_2 of category ψ_2 with 10 m distance, and p_3 of category ψ_2 with 80 m distance. The resulting vectors (assuming there are only three categories) are $[0, 2, 1]$ and $[200, 10, 50]$.

- Space2vec: In contrast to hand-crafted features based on distance and category counts, there is the option to *learn* coordinate representations. The task of finding an efficient and informative representation of points, dependent on their coordinates and POI context, was tackled recently in work on space embeddings. We employ the state-of-the-art *space-to-vec* approach by Mai et al. [52]. Inspired by word embeddings in natural language processing, the idea is to learn a compact vector representation for points. The training is based on a supervised learning task, namely to distinguish surrounding points from unrelated, arbitrary distant samples that were drawn as negative samples. We deploy their public code base¹⁰ to train the algorithm on our POI datasets \mathcal{P} , including the first *two* category levels. Specifically, we split \mathcal{P} into training, validation (10%), and testing set (10%) and employ the *joined* approach by Mai et al. [52]; i.e., training a location decoder and a spatial context decoder jointly. We set the embedding size to 16 but retained all other parameters as suggested by the authors. The model, which was trained only on \mathcal{P} , can be applied

⁹ For brevity, we omit l_i^u here.

¹⁰ <https://github.com/gengchenmai/space2vec>.

on a new location given its coordinates and its spatial context (coordinates and categories of the surrounding POIs) as input.

Machine learning model

We chose the XGB approach over other machine learning models for its interpretability and its suitability for unbalanced data, rendering it superior in many applications. Nevertheless, we also implemented a multi-layer perceptron (MLP) for comparison. A model was implemented with two layers of 128 neurons respectively, with dropout regularization, ReLU activation and a softmax function in the output layer. The network was trained with the AdamOptimizer (learning rate 0.001) and with early stopping. For the XGBoost model, we utilize the XGBoost implementation in the `xgboost` Python package¹¹ and only tune the parameter that determines the maximum depth of the base learners. A depth of 10 turned out most suitable in our experiments. The MLP also exhibits good place categorization ability, but was consistently inferior to XGB. For example, with the Foursquare data for NYC and an obfuscation radius of 100 m, the accuracy is 52.2% for the MLP compared to 59.4% for XGB (41.6% vs 49.8% for 200 m obfuscation, etc.). We, therefore, only report the results for XGB in this study.

Location masking

A simple protection method for the use of location-based services is a random displacement of the coordinates to mask the real location. For example, iPhone users can withhold the precise locations from applications and only allow them to access the “approximate” location. Here, we utilize location obfuscation to model imprecise GNSS data or basic data protection. The user’s location is simply replaced by a new location sampled from a uniform distribution within a given radius r (see Fig. 1a). Note that we focus on the obfuscation of the spatial information and leave the possibility of masking temporal information as in [59] for future work on semantic privacy. After the location masking step (Fig. 1a), the raw (and obfuscated) spatio-temporal data are featurized (Fig. 1b) by deriving temporal features from the check-in time and spatial features from the coordinates matched with public POI data.

Data split

We test for the attacker’s accuracy by splitting the data into train and test sets, as shown in Fig. 1c. By default, the dataset is split by user, i.e., 10% of the users are taken as the test set while the model is trained on 90%. In practice, we report all results upon *tenfold cross validation* such that all users were part of the test set once. The results simulate the scenario where the attacker obtains a labeled train dataset from a specific region and utilizes it to train an ML model with the goal to infer location profiles of new users but in the same region. However, the attacker may not always have labeled data from exactly the same spatial region. To analyze this scenario, we additionally simulate the attack with a *spatial split*. In detail, the dataset is divided by separating the x- and y-coordinates in a 3×3 grid to yield nine roughly equal-sized subsets. The samples from each grid cell are used as the test set once.

¹¹ https://xgboost.readthedocs.io/en/stable/python/python_intro.html.

Appendix

Specific scenarios of semantic privacy attacks

To further elaborate on potential risks from semantic privacy attacks and to counter the “I have nothing to hide” argument [79], we list specific scenarios in the following. These scenarios assume that a data broker has obtained location data of a user.

1. The data indicates that a user spends much time in recreational areas and regularly visits outdoor shops. The user is selected for targeted ads about travelling and hiking.
2. A user was detected to visit a pharmacy or medical care unit frequently. This information is sold to an insurance company that adapts their policy offers accordingly, or targets only users without such behaviour in their ad campaign.
3. A user is observed to visit church regularly. The election campaign of the conservative party targets this person by providing specific information on the religious values of the party.
4. A user known to be a young woman is detected in a nightclub and outside of her home location at night, and is found to visit a pharmacy in the next morning. A data broker may infer that she bought a contraceptive pill.
5. The location data reveals a user to be increasingly active at nights visiting bars and clubs. This person may be targeted for drug testing after data breaches.

Scenario comparison results

Table 1 provides an overview of the considered scenarios and metrics, and lists the results for protected location data with an obfuscation radius of 100 m as an example. In the first scenario, both the check-in train data and the POI data are taken from the Foursquare dataset, and a machine-learning attack with the XGBoost model trained on spatio-temporal features can achieve a place categorization accuracy of 61.6%. This is remarkable considering that there are 12 categories. It further translates to a user profiling error¹² of 0.124 which corresponds to a top-5 user re-identification accuracy (hit@5) of 40.4%. In other words, given (obfuscated) location data and temporal visitation patterns, the attacker can infer an approximate profile $\hat{p}(u)$ for user u such that, in 40.4% cases, his real profile $p(u)$ is among the five real user profiles most similar to $\hat{p}(u)$.

Furthermore, as shown in the confusion matrix in Fig. 3, the place classification accuracy clearly differs between categories. A more detailed analysis of this effect is provided in Fig. 11. At stronger obfuscation, many places are erroneously labeled as “Dining” or “Travel and Transportation”. Therefore, the sensitivity for these places remains high, while decreasing for the other place categories.

Tokyo vs NYC—effect of location obfuscation

In Fig. 12, we compare the results on two diverse cities, NYC and Tokyo. While the results show the same decrease- and convergence behavior in both cities, the accuracy is generally larger for Tokyo. This is, however, due to a stronger label imbalance for the

¹² The user profiling error is the average Euclidean distance between real and predicted profiles, $E_{\hat{p}(u), p(u)}$

Table 1 Comparing attack-scenarios. For the place categorization task, the accuracy over location-visit events is reported. User-profiling error, user identification accuracy and privacy loss measure the success of the attacker in user-profiling. The table only shows the results for an obfuscation radius of $r = 100$

Check-in data	POI data	Method	Place categorization	User profiling				
			Accuracy	Profiling error	Top-5 identification accuracy	Privacy loss (median)		
Foursquare (NYC and Tokyo)	Foursquare	Random	0.159	0.235	0.003	1.000		
		XGB (temporal)	0.344	0.205	0.023	1.112		
		spatial join	0.459	0.208	0.183	2.296		
		XGB (spatial)	0.580	0.140	0.328	5.574		
	OSM	XGB (spatio-temporal)	0.616	0.124	0.404	11.001		
		Random	0.157	0.235	0.002	1.000		
		Spatial join	0.224	0.341	0.013	1.051		
		XGB (temporal)	0.328	0.206	0.022	1.116		
		XGB (spatial)	0.476	0.172	0.117	1.949		
		XGB (spatio-temporal)	0.503	0.157	0.160	2.316		
		Yumuv study (Switzerland)	Foursquare	Random	0.370	0.367	0.010	1.000
				Spatial join	0.453	0.698	0.009	1.000
XGB (spatial)	0.541			0.331	0.034	1.095		
XGB (temporal)	0.553			0.333	0.038	1.108		
XGB (spatio-temporal)	0.618			0.293	0.065	1.225		
OSM	Random		0.372	0.367	0.010	1.000		
	Spatial join		0.384	0.698	0.009	1.000		
	XGB (spatial)		0.541	0.332	0.034	1.093		
	XGB (temporal)		0.553	0.333	0.038	1.108		
	XGB (spatio-temporal)		0.623	0.289	0.060	1.256		

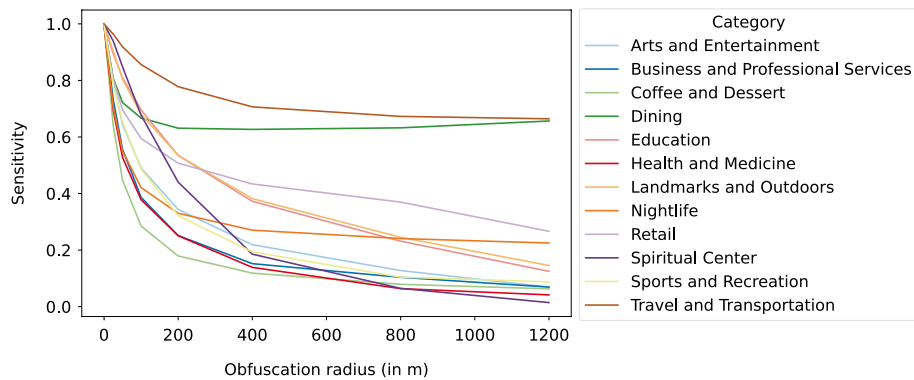


Fig. 11 Sensitivity of place recognition by category




Fig. 12 Place categorization accuracy of Tokyo compared to NYC

places in Tokyo which leads to a better performance of the random baseline (grey lines Fig. 12).

Effect of POI data completeness on place categorization task

Figure 13 provides the results for the first task (place categorization) by POI data completeness, complementing the user-identification results in Fig. 9. As was already seen for user-identification, the accuracy decreases disproportionately.

User profiling error

Corresponding to the user identification accuracy presented in Fig. 6, we show the changes of the profiling errors in Fig. 14, where the error is the Euclidean distance between the user’s real profile vector $p_c(u)$ and the estimated profile vector $\hat{p}_c(u)$.




Fig. 13 Effect of reduced POI context data quality on task 1 (place categorization)




Fig. 14 User profiling error (Euclidean distance between true and predicted user profile) by obfuscation radius

Study on GNSS tracking data

The Foursquare check-in dataset is suitable for our analysis since it is reliably labeled with detailed place categories. For comparison, we evaluate results on a GNSS-based tracking dataset from the so-called yumuv study. yumuv is a tracking study in Switzerland that was conducted to investigate the mobility behavior with a micro-mobility-bundle subscription [71]. Participants of the study were tracked via the Myway app¹³ for 2–3 months and manually labeled their activities. The app provides a preprocessed version of the raw GNSS data, namely (labeled) staypoints and triplegs. We further use the Python library Trackintel [55] to group staypoints into *locations* via DBSCAN. As for Foursquare, we remove the “home” check-ins as well as further categories that are uninformative, such as “wait”, “errand”, and “unknown”. This leads to a set of five location categories, namely $C = \{\text{Work}, \text{Education}, \text{Sports and Recreation}, \text{Dining}, \text{Leisure}\}$. The categorical distribution is shown in Fig. 16b.

The dataset is more challenging due to noisy GNSS tracking data as well as unreliable and incomplete activity labels (the labels were estimated by an app and manually corrected by the user). Additionally, different users hardly visit the same places, and, thus, a training dataset has limited value for an attacker. Indeed, the main source of information for the attacker are temporal features in the case of the yumuv data. Figure 15 shows that the attacker can still achieve semantic inference that is significantly better than random; however, spatial information only plays a limited role. Consequently, location protection hardly hinders the semantic attack, and the performance converges already at an obfuscation radius of 100 m. In Table 1, other metrics are listed, and it is also shown that the danger for privacy on a user level is limited, with a top-5 user identification accuracy of only up to 7.7% (100 m location obfuscation). We hypothesize that the main reason for the limited value of spatial context data for the attacker lies in the activity categories of

¹³ <https://www.sbb.ch/en/timetable/mobile-apps/myway.html>.

the yumuv study (see Fig. 16b), including “Work” and “Leisure” as the largest classes, which are too generic and too difficult to locate with POI data.

Dataset analysis

In Fig. 16, the distribution of category types over visited places is shown for Foursquare (Fig. 16a) and the yumuv (Fig. 16b) data.

Analyzing place category autocorrelation in a semivariogram

The categories of user locations are predicted based on information from surrounding public POIs. Even if the location is obfuscated, the category of the closest POI is a good predictor for the location’s category. This is due to the spatial autocorrelation of place categories. For example, there are urban areas with many restaurants in one street, such that there is a high probability of finding several nearby places with the same category. We relate our results to the spatial autocorrelation of POIs by comparing our place categorization accuracy (Fig. 2) to the semivariogram of place category differences. To compute the variogram, we bin the distances and report the category correspondence

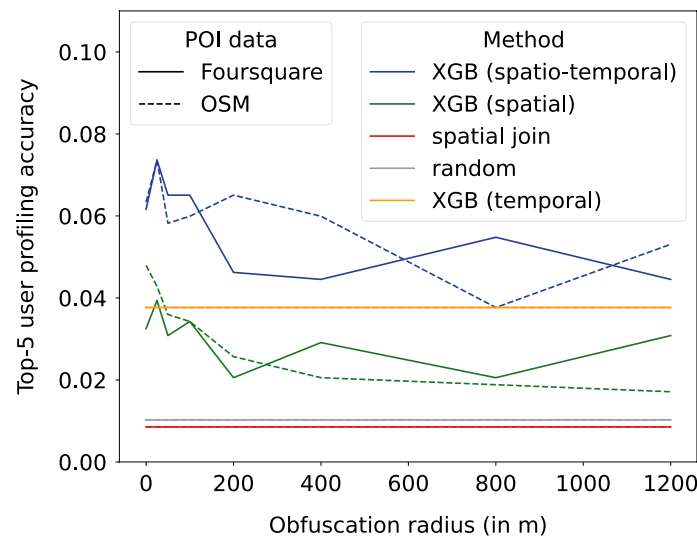


Fig. 15 Profiling results on GNSS tracking data. The profiling accuracy is low even with low location obfuscation, and the relation between obfuscation radius and profiling accuracy is noisy




Fig. 16 Location category counts in Foursquare (a) and Yumuv (b) datasets




Fig. 17 Semivariogram for measuring spatial autocorrelation of place categories. The y-axis shows the percentage of POI-pairs with different categories

in each bin, thereby accounting for the categorical nature of place categories and the continuous nature of the place coordinates. In detail, given a minimum and maximum distance d_{min}, d_{max} , the category variance is computed as

$$\gamma(d_{min}, d_{max}) = \frac{\sum_{l_1, l_2 \in N(d_{min}, d_{max})} \mathbb{1}[c(l_1) == c(l_2)]}{|N(d_{min}, d_{max})|}$$

where $N(d_{min}, d_{max})$ is the set of all pairs of locations with $d_{min} < d(l_1, l_2) \leq d_{max}$, and $\mathbb{1}$ is the indicator function that counts the number of pairs with corresponding category. Due to the high number of POIs in the dataset, we only take a 20×20 km subregion of New York City and sample pairs of POIs randomly.

Figure 17 shows the change of γ with increasing d_{min} and d_{max} , as common in a semivariogram. Even places with less than 25 m distance only have a 33% chance of their categories to correspond ($\gamma(0, 25) = 0.67$). At a distance of more than 3.2km, the category variance converges to around 89%, which is slightly lower than the expected variance of 8.3% for 12 categories. However, the imbalance of categories explains this difference. The semivariogram confirms our finding in Fig. 2 that spatial context data is hardly informative for location data that is obfuscated more than 1 km.

Experimental results for six place categories

The choice of place categories substantially effects the attacker's success. There is a trade-off between the place labelling accuracy and the utility of the user profiles: The more categories, the more fine-grained and informative the user profiles, but the less reliable the classification. To demonstrate, we construct a new set of six categories by merging the 12 original place categories by their semantic similarity, yielding the set {"Sports, Landmarks, Outdoors", "Nightlife, Arts, Entertainment", "Dining, Coffee, Desert", "Retail, Business, Professional Services", "Travel and Transportation", "Education, Spiritual Center, Health and Medicine"}. Figure 18 shows the place labelling accuracy and top-5 user re-identification accuracy with these six instead of twelve categories (only NYC, Foursquare check-ins and POI data, user split). The place labelling accuracy is higher than with 12 categories, whereas the user re-identification performance decreases

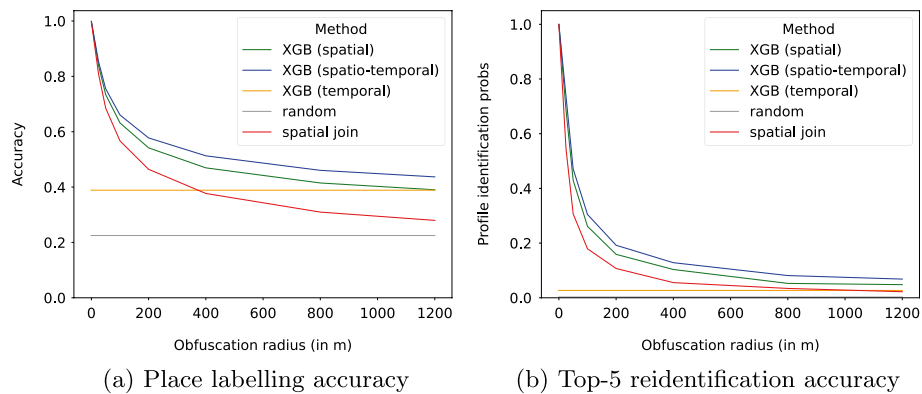


Fig. 18 Experimental results with lower granularity of activity distinction (6 categories). As expected, the place labelling accuracy is higher than with 12 categories, converging to 40% **(a)**. However, a lower granularity of the categories results in less expressive user profiles, which is reflected in a lower probability to re-identify users by their behavioural profile **(b)**

due to the lower granularity of the user profiles. For example, at an obfuscation radius of 200 m, the categorization is correct in 58% cases (compared to 50% cases with 12 categories), but the top-5 user re-identification accuracy is as low as 19% (compared to 35% with 12 categories). From an attacker's perspective, it is interesting to explore which set of categories provides the most informative, yet predictable user profiles.

Abbreviations

ML	Machine learning
XGB	XGBoost algorithm
POI	Point of interest
TF-IDF	Term Frequency-Inverse Document Frequency
OSM	Open Street Map
GNSS	Global navigation satellite system
TTP	Trusted third-party

Acknowledgements

Not applicable.

Author contributions

NW, OK and KJ conceptualized the project. NW and OK performed the literature research. NW developed the methodology, implemented the algorithms, prepared all visualizations and wrote the main manuscript draft. KJ revised the manuscript. OK and MR supervised the project and reviewed the manuscript.

Funding

Open access funding provided by Swiss Federal Institute of Technology Zurich.

Availability of data and materials

The Foursquare data is publicly available at <https://sites.google.com/site/yangdingqi/home/foursquare-dataset?pli=1>. All source code for reproducing our results is published on GitHub: https://github.com/mie-lab/trip_purpose_privacy.

Declarations

Ethics approval and consent to participate

Not applicable.

Consent for publication

Not applicable.

Competing interests

The authors declare that they have no competing interests.

Received: 22 May 2023 Accepted: 28 January 2024

Published online: 12 March 2024

References

1. Al Hasan Haldar N, Li J, Reynolds M, Sellis T, Yu JX. Location prediction in large-scale social networks: an in-depth benchmarking study. *VLDB J*. 2019;28(5):623–48.
2. Alrayes F, Abdelmoty A. No place to hide: a study of privacy concerns due to location sharing on geo-social networks. *Int J Inf Secur*. 2014;7(3/4):62–75.
3. An N, Chen M, Lian L, Li P, Zhang K, Yu X, Yin Y. Enabling the interpretability of pretrained venue representations using semantic categories. *Knowl-Based Syst*. 2022;235:107623.
4. Andrés ME, Bordenabe NE, Chatzikokolakis K, Palamidessi C. Geo-indistinguishability: differential privacy for location-based systems. In: *Proceedings of the 2013 ACM SIGSAC conference on computer & communications security*. 2013. p. 901–14.
5. Banerjee S. Geosurveillance, location privacy, and personalization. *J Public Policy Mark*. 2019;38(4):484–99.
6. Bao J, Zheng Y, Wilkie D, Mokbel M. Recommendations in location-based social networks: a survey. *Geoinformatica*. 2015;19(3):525–65.
7. Barth S, De Jong MD. The privacy paradox—investigating discrepancies between expressed privacy concerns and actual online behavior—a systematic literature review. *Telemat Inform*. 2017;34(7):1038–58.
8. Brockmann D, Hufnagel L, Geisel T. The scaling laws of human travel. *Nature*. 2006;439(7075):462–5.
9. Cerf S, Primault V, Boutet A, Mokhtar SB, Birke R, Bouchenak S, Chen LY, Marchand N, Robu B. Pulp: achieving privacy and utility trade-off in user mobility data. In: *2017 IEEE 36th symposium on reliable distributed systems (SRDS)*. 2017. p. 164–73.
10. Charleux L, Schofield K. True spatial k-anonymity: Adaptive areal elimination vs. adaptive areal masking. *Cartogr Geogr Inf Sci*. 2020;47(6):537–49.
11. Chen C, Ma J, Susilo Y, Liu Y, Wang M. The promises of big data and small data for travel behavior (aka human mobility) analysis. *Transp Res Part C Emerg Technol*. 2016;68:285–99.
12. Chen Q, Poorthuis A. Identifying home locations in human mobility data: an open-source r package for comparison and reproducibility. *Int J Geogr Inf Sci*. 2021;35(7):1425–48.
13. Chen T, Guestrin C. Xgboost: a scalable tree boosting system. In: *Proceedings of the 22nd ACM SIGKDD international conference on knowledge discovery and data mining*. 2016. p. 785–94.
14. Cheng W, Wen R, Huang H, Miao W, Wang C. OPTDP: towards optimal personalized trajectory differential privacy for trajectory data publishing. *Neurocomputing*. 2022;472:201–11.
15. Crain M. The limits of transparency: data brokers and commodification. *New Media Soc*. 2018;20(1):88–104.
16. Crandall DJ, Backstrom L, Cosley D, Suri S, Huttenlocher D, Kleinberg J. Inferring social ties from geographic coincidences. *Proc Natl Acad Sci*. 2010;107(52):22436–41.
17. Cui Y, Meng C, He Q, Gao J. Forecasting current and next trip purpose with social media data and Google Places. *Transp Res Part C Emerg Technol*. 2018;97:159–74.
18. de Montjoye Y-A, Hidalgo CA, Verleysen M, Blondel VD. Unique in the crowd: the privacy bounds of human mobility. *Sci Rep*. 2013;3(1):1376.
19. Do TMT, Gatica-Perez D. The places of our lives: visiting patterns and automatic labeling from longitudinal smartphone data. *IEEE Trans Mobile Comput*. 2014;13(3):638–48.
20. Du X, Zhu H, Zheng Y, Lu R, Wang F, Li H. A semantic-preserving scheme to trajectory synthesis using differential privacy. *IEEE Internet Things J*. 2023;10(5):13784–97.
21. Duckham M, Kulik L. A formal model of obfuscation and negotiation for location privacy. In: Hutchison D, Kanade T, Kittler J, Kleinberg JM, Mattern F, Mitchell JC, Naor M, Nierstrasz O, Pandu Rangan C, Steffen B, Sudan M, Terzopoulos D, Tygar D, Vardi MY, Weikum G, Gellersen HW, Want R, Schmidt A, editors. *Pervasive computing*, vol. 3468. Berlin: Springer; 2005. p. 152–70.
22. Duckham M, Kulik L. Location privacy and location-aware computing. In: *Dynamic and mobile GIS*. Boca Raton: CRC Press; 2006. p. 63–80.
23. Dwork C. Differential privacy: a survey of results. In: *International conference on theory and applications of models of computation*. Springer; 2008. p. 1–19.
24. Efstathiades H, Antoniadis D, Pallis G, Dikaiakos MD. Identification of key locations based on online social network activity. In: *2015 IEEE/ACM international conference on advances in social networks analysis and mining (ASONAM)*. IEEE; 2015. p. 218–25.
25. Falcone D, Mascolo C, Comito C, Talia D, Crowcroft J. What is this place? Inferring place categories through user patterns identification in geo-tagged tweets. In: *6th international conference on mobile computing, applications and services*. 2014. p. 10–9.
26. Gao S, Janowicz K, Couclelis H. Extracting urban functional regions from points of interest and human activities on location-based social networks. *Trans GIS*. 2017;21(3):446–67.
27. Gao S, Rao J, Liu X, Kang Y, Huang Q, App J. Exploring the effectiveness of geomasking techniques for protecting the geoprivacy of twitter users. *J Spat Inf Sci*. 2019;19:105–29.
28. Gao X, Firner B, Sugrim S, Kaiser-Pendergrast V, Yang Y, Lindqvist J. Elastic pathing: your speed is enough to track you. In: *Proceedings of the 2014 ACM international joint conference on pervasive and ubiquitous computing*. 2014. p. 975–86.
29. Georgiadou Y, de By RA, Kounadi O. Location privacy in the wake of the GDPR. *ISPRS Int J Geo-Inf*. 2019;8(3):157.
30. Golle P, Partridge K. On the anonymity of home/work location pairs. In: *Pervasive computing: 7th international conference, pervasive 2009, Nara, Japan, May 11–14, 2009. Proceedings 7*. Springer; 2009. p. 390–7.
31. Götz M, Nath S, Gehrke J. Maskit: privately releasing user context streams for personalized mobile applications. In: *Proceedings of the 2012 ACM SIGMOD international conference on management of data*. 2012. p. 289–300.
32. Grinsztajn L, Oyallon E, Varoquaux G. Why do tree-based models still outperform deep learning on typical tabular data? In: *Thirty-sixth conference on neural information processing systems datasets and benchmarks track*. 2022.
33. Gruteser M, Grunwald D. Anonymous usage of location-based services through spatial and temporal cloaking. In: *Proceedings of the 1st international conference on mobile systems, applications and services*. 2003. p. 31–42.

34. Guo C, Pleiss G, Sun Y, Weinberger KQ. On calibration of modern neural networks. In: International conference on machine learning. PMLR; 2017. p. 1321–30.
35. Gurung S, Lin D, Jiang W, Hurson A, Zhang R. Traffic information publication with privacy preservation. *ACM Trans Intell Syst Technol.* 2014;5(3):1–26.
36. Han J, Owusu E, Nguyen LT, Perrig A, Zhang J. Accomplice: location inference using accelerometers on smartphones. In 2012 fourth international conference on communication systems and networks (COMSNETS 2012). IEEE; 2012. p. 1–9.
37. Haydari A, Zhang M, Chuah C-N, Macfarlane J, Peisert S. Adaptive differential privacy mechanism for aggregated mobility dataset. *arXiv preprint.* 2021. [arXiv:2112.08487](https://arxiv.org/abs/2112.08487)
38. Huang H, Gartner G, Krisp JM, Raubal M, Van de Weghe N. Location based services: ongoing evolution and research agenda. *J Location Based Serv.* 2018;12(2):63–93.
39. Huang S, Papernot N, Goodfellow I, Duan Y, Abbeel P. Adversarial attacks on neural network policies. *arXiv preprint.* 2017. [arXiv:1702.02284](https://arxiv.org/abs/1702.02284).
40. Jain P, Gyanchandani M, Khare N. Differential privacy: its technological prescriptive using big data. *J Big Data.* 2018;5(1):1–24.
41. Janowicz K. Observation-driven geo-ontology engineering. *Trans GIS.* 2012;16(3):351–74.
42. Jenson S, Reeves M, Tomasini M, Menezes R. Mining location information from users' spatio-temporal data. In: 2017 IEEE smartworld, ubiquitous intelligence & computing, advanced & trusted computed, scalable computing & communications, cloud & big data computing, internet of people and smart city innovation. 2017. p. 1–7.
43. Jiang H, Li J, Zhao P, Zeng F, Xiao Z, Iyengar A. Location privacy-preserving mechanisms in location-based services: a comprehensive survey. *ACM Comput Surv.* 2021;54(1):1–36.
44. Keßler C, McKenzie G. A geoprivacy manifesto. *Trans GIS.* 2018;22(1):3–19.
45. Khan SI, Khan ABA, Hoque ASML. Privacy preserved incremental record linkage. *J Big Data.* 2022;9(1):1–27.
46. Kounadi O, Lampoltshammer TJ, Leitner M, Heistracher T. Accuracy and privacy aspects in free online reverse geocoding services. *Cartogr Geogr Inf Sci.* 2013;40(2):140–53.
47. Kounadi O, Resch B, Petutschnig A. Privacy threats and protection recommendations for the use of geosocial network data in research. *Soc Sci.* 2018;7(10):191.
48. Krumm J. Inference attacks on location tracks. In: Proceedings of the 5th international conference on Pervasive computing, PERVASIVE'07. Berlin: Springer-Verlag; 2007. p. 127–43.
49. Li H, Zhu H, Du S, Liang X, Shen X. Privacy leakage of location sharing in mobile social networks: attacks and defense. *IEEE Trans Dependable Secure Comput.* 2016;15(4):646–60.
50. Liu B, Ding M, Shaham S, Rahayu W, Farokhi F, Lin Z. When machine learning meets privacy: a survey and outlook. *ACM Comput Surv.* 2021;54(2):1–36.
51. Liu R, Buccapatnam S, Gifford WM, Sheopuri A. An unsupervised collaborative approach to identifying home and work locations. In: 2016 17th IEEE international conference on mobile data management (MDM), vol. 1. 2016. p. 310–7.
52. Mai G, Janowicz K, Yan B, Zhu R, Cai L, Lao N. Multi-scale representation learning for spatial feature distributions using grid cells. *arXiv preprint.* 2020. [arXiv:2003.00824](https://arxiv.org/abs/2003.00824).
53. Manousakas D, Mascolo C, Beresford AR, Chan D, Sharma N. Quantifying privacy loss of human mobility graph topology. *Proc Privacy Enhancing Technol.* 2018;2018(3):5–21.
54. Martin H, Wiedemann N, Suel E, Hong Y, Xin Y. Influence of tracking duration on the privacy of individual mobility graphs. In: Proceedings of the 17th international conference on location-based services. Technical University of Munich; 2022.
55. Martin H, Hong Y, Wiedemann N, Bucher D, Raubal M. Trackintel: an open-source python library for human mobility analysis. *Comput Environ Urban Syst.* 2023;101: 101938.
56. McKenzie G, Janowicz K. Where is also about time: a location-distortion model to improve reverse geocoding using behavior-driven temporal semantic signatures. *Comput Environ Urban Syst.* 2015;54:1–13.
57. McKenzie G, Zhang H. Platial k-anonymity: improving location anonymity through temporal popularity signatures. In: 12th International Conference on Geographic Information Science (GIScience 2023). Schloss Dagstuhl-Leibniz-Zentrum für Informatik; 2023.
58. McKenzie G, Janowicz K, Gao S, Yang J-A, Hu Y. Poi pulse: a multi-granular, semantic signature-based information observatory for the interactive visualization of big geosocial data. *Cartogr Int J Geogr Inf Geovis.* 2015;50(2):71–85.
59. McKenzie G, Janowicz K, Seidl D. Geo-privacy beyond coordinates. In: Geospatial data in a changing world. Cham: Springer; 2016. p. 157–75.
60. McKenzie G, Romm D, Zhang H, Brunila M. PrivyTo: a privacy preserving location sharing platform. *Trans GIS.* 2022;26:16.
61. Miranda-Pascual À, Guerra-Balboa P, Parra-Arnau J, Forné J, Strufe T. SoK: differentially private publication of trajectory data. *Proc Priv Enhancing Technol.* 2023;2:496–516.
62. Montini L, Rieser-Schüssler N, Horni A, Axhausen KW. Trip purpose identification from GPS tracks. *Transp Res Rec.* 2014;2405(1):16–23.
63. Nelson T, Goodchild M, Wright D. Accelerating ethics, empathy, and equity in geographic information science. *Proc Natl Acad Sci.* 2022;119(19): e2119967119.
64. Olteanu A-M, Huguenin K, Shokri R, Humbert M, Hubaux J-P. Quantifying interdependent privacy risks with location data. *IEEE Trans Mobile Comput.* 2016;16(3):829–42.
65. Pei J, Xu J, Wang Z, Wang W, Wang K. Maintaining k-anonymity against incremental updates. In: 19th international conference on scientific and statistical database management (SSDBM 2007). IEEE; 2007. p. 5–5.
66. Penha Natal ID, Avellar Campos Cordeiro RD, Garcia ACB. Activity recognition model based on GPS data, points of interest and user profile. In: International symposium on methodologies for intelligent systems. Springer; 2017. p. 358–67.
67. Pontes T, Vasconcelos M, Almeida J, Kumaraguru P, Almeida V. We know where you live: privacy characterization of foursquare behavior. In: Proceedings of the 2012 ACM conference on ubiquitous computing. 2012. p. 898–905.

68. Pötzsch S. Privacy awareness: a means to solve the privacy paradox? In: IFIP summer school on the future of identity in the information society. Springer; 2008. p. 226–36.
69. Qiu G, Tang G, Li C, Guo D, Shen Y, Gan Y. Behavioral-semantic privacy protection for continual social mobility in mobile-internet services. *IEEE Internet Things J.* 2023;11(1):462–77.
70. Ram Mohan Rao P, Murali Krishna S, Siva Kumar A. Privacy preservation techniques in big data analytics: a survey. *J Big Data.* 2018;5:1–12.
71. Reck DJ, Martin H, Axhausen KW. Mode choice, substitution patterns and environmental impacts of shared and personal micro-mobility. *Transp Res Part D Transp Environ.* 2022;102: 103134.
72. Rhee I, Shin M, Hong S, Lee K, Kim SJ, Chong S. On the levy-walk nature of human mobility. *IEEE/ACM Trans Netw.* 2011;19(3):630–43.
73. Rossi L, Walker J, Musolesi M. Spatio-temporal techniques for user identification by means of GPS mobility data. *EPJ Data Sci.* 2015;4(1):11.
74. Seidl DE, Jankowski P, Tsou M-H. Privacy and spatial pattern preservation in masked GPS trajectory data. *Int J Geogr Inf Sci.* 2016;30(4):785–800.
75. Shen L, Stopher PR. A process for trip purpose imputation from global positioning system data. *Transp Res Part C Emerg Technol.* 2013;36:261–7.
76. Sherman J. Data brokers and sensitive data on us individuals. Duke University Sanford Cyber Policy Program. 2021;9.
77. Shokri R. Quantifying and protecting location privacy. *IT-Inf Technol.* 2015;57(4):257–63.
78. Shokri R, Theodorakopoulos G, Le Boudec JY, Hubaux J-P. Quantifying location privacy. In: 2011 IEEE symposium on security and privacy. 2011. p. 247–62.
79. Solove DJ. I've got nothing to hide and other misunderstandings of privacy. *San Diego L Rev.* 2007;44:745.
80. Sreekumar S, Gündüz D. Optimal privacy-utility trade-off under a rate constraint. In: 2019 IEEE international symposium on information theory (ISIT). IEEE; 2019. p. 2159–63.
81. Sweeney L. k-anonymity: a model for protecting privacy. *Int J Uncertain Fuzziness Knowl-Based Syst.* 2002;10(05):557–70.
82. Tenkanen H. pyrosm v0.6.1. 2022.
83. Tu Z, Zhao K, Xu F, Li Y, Su L, Jin D. Protecting trajectory from semantic attack considering k-anonymity, l-diversity, and t-closeness. *IEEE Trans Netw Serv Manag.* 2019;16(1):264–78.
84. Wernke M, Skvortsov P, Dürr F, Rothermel K. A classification of location privacy attacks and approaches. *Pers Ubiquit Comput.* 2014;18:163–75.
85. Xiao G, Juan Z, Zhang C. Detecting trip purposes from smartphone-based travel surveys with artificial neural networks and particle swarm optimization. *Transp Res Part C Emerg Technol.* 2016;71:447–63.
86. Xie M, Yin H, Wang H, Xu F, Chen W, Wang S. Learning graph-based poi embedding for location-based recommendation. In: Proceedings of the 25th ACM international on conference on information and knowledge management. 2016. p. 15–24.
87. Yan B, Janowicz K, Mai G, Gao S. From itdl to place2vec: reasoning about place type similarity and relatedness by learning embeddings from augmented spatial contexts. In: Proceedings of the 25th ACM SIGSPATIAL international conference on advances in geographic information systems. 2017. p. 1–10.
88. Yan Y, Xu F, Mahmood A, Dong Z, Sheng QZ. Perturb and optimize users' location privacy using geo-indistinguishability and location semantics. *Sci Rep.* 2022;12(1):1–20.
89. Yang D, Zhang D, Zheng VW, Yu Z. Modeling user activity preference by leveraging user spatial temporal characteristics in LBSNS. *IEEE Trans Syst Man Cybern Syst.* 2015;45(1):129–42.
90. Yang D, Zhang D, Qu B, Cudré-Mauroux P. Privcheck: privacy-preserving check-in data publishing for personalized location based services. In: Proceedings of the 2016 ACM international joint conference on pervasive and ubiquitous computing. 2016. p. 545–56.
91. Ye M, Shou D, Lee W-C, Yin P, Janowicz K. On the semantic annotation of places in location-based social networks. In: Proceedings of the 17th ACM SIGKDD international conference on knowledge discovery and data mining—KDD '11. San Diego: ACM Press; 2011. p. 520.
92. Ye M, Yin P, Lee W-C, Lee D-L. Exploiting geographical influence for collaborative point-of-interest recommendation. In: Proceedings of the 34th international ACM SIGIR conference on research and development in information retrieval. 2011. p. 325–34.
93. Ying JJ-C, Lu EH-C, Lee W-C, Weng T-C, Tseng VS. Mining user similarity from semantic trajectories. In: Proceedings of the 2nd ACM SIGSPATIAL international workshop on location based social networks. 2010. p. 19–26.
94. Yuan Y, Raubal M. Analyzing the distribution of human activity space from mobile phone usage: an individual and urban-oriented study. *Int J Geogr Inf Sci.* 2016;30(8):1594–621.
95. Zhang L, Pfoser D. Using OpenStreetMap point-of-interest data to model urban change—a feasibility study. *PLoS ONE.* 2019;14(2): e0212606.
96. Zhang X, Huang H, Huang S, Chen Q, Ju T, Du X. A context-aware location differential perturbation scheme for privacy-aware users in mobile environment. *Wirel Commun Mobile Comput.* 2018.
97. Zhu D, Zhang F, Wang S, Wang Y, Cheng X, Huang Z, Liu Y. Understanding place characteristics in geographic contexts through graph convolutional neural networks. *Ann Am Assoc Geogr.* 2020;110(2):408–20.

Publisher's Note

Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.