Informative Yet Unrevealing: Semantic Obfuscation for Location Based Services

Jing Yang, Zack Zhu, Julia Seiter, Gerhard Tröster
Wearable Computing Lab
ETH Zurich
yangj@ethz.ch, {zack.zhu, julia.seiter, troester}@ife.ee.ethz.ch

ABSTRACT

The preservation of geo-privacy is a critical consideration for location-based service (LBS) providers. Unfortunately, a trade-off typically exists between the quality of location-based services and revealing of private information (e.g. geo-coordinates) to obtain such services. In this work, we develop semantic obfuscation methods, which allow a trusted third-party to convert revealing geo-coordinates into highly anonymous semantic features. Following, LBS providers can operate directly via location semantics to deliver the necessary services. Using a large-scale travel survey dataset, we evaluate our obfuscation approach while considering a common user-intention prediction problem. Our results demonstrate that our approach is capable of significantly obfuscating user location while maintaining LBS quality. On average, we show that the k-anonymity measure increases by 15.22 times while the quality of prediction drops only 3.24%.

Categories and Subject Descriptors

H.2.8 [**Database Applications**]: Spatial Databases and GIS, Data Mining

General Terms

Theory

Keywords

Location Based Services, Geo-Privacy, Semantic Obfuscation, Obfuscation Metrics

1. INTRODUCTION

The ubiquitous presence of GPS-enabled mobile phones have facilitated a vast number of location-based services (LBS). Such services typically require the specific geo-coordindates of user location, which can then be used to provide just-in-time and context-relevant services. Such services could be, for example, providing the nearest restaurant information to users [16], showing the nearby gas stations to drivers [9], notifying customers of special promotion in stores [16], etc. As LBS depend on privacy-sensitive real-time geo-coordinates, the risk of malicious parties intruding on people's privacy is a critical concern to users of LBS. On the other hand, the precision of geo-coordinates is critical for understanding user context and providing high-quality LBS [20]. As such, there is often a trade-off between geo-privacy preservation and LBS quality.

In our work, we aim to perturb geo-information via location obfuscation approaches to protect user privacy *while* maintaining the quality of service of LBS. To this end, we introduce a trusted third-party that converts absolute geo-data (e.g. GPS coordinates) into relative geo-features (semantics of nearby venues, e.g. office, hospital, elementary school, etc.), which is typically conducted by LBS providers anyway as a feature extraction process. Given relative geo-features, we propose semantic obfuscation methods to semantically obfuscate and modify the feature vector. To preserve the quality-of-service, we obfuscate with consideration of the final provided service. As a case study, we evaluate our obfuscation methods on a travel purpose prediction by [20]. There, the LBS (prediction of travel purposes) is provided using relative geo-features, however, without any obfuscation. Using their approach as a benchmark, we show that malicious parties can easily reverse-engineer for absolute user locations from relative geo-features. Using our proposed obfuscation method, preservation of geo-privacy significantly increases while LBS quality experience marginal decrease.

Concretely, our contributions beyond the state-of-the-art work are:

- An effective semantic-level obfuscation method to preserve geo-privacy while maintaining LBS quality.

- Simple and intuitive obfuscation metrics to quantify the privacy preservation level.

- Evaluation of the methods on large-scale human mobility data to substantiate the usefulness of our methods.

In the rest of this paper, we first review existing work in Section 2. Details of our methodologies are presented in Section 3. We evaluate our methods and present experimental results in Section 4 and further discussion is provided in Section 5. Finally we conclude the paper and outline future work in Section 6.

GeoPrivacy'15, November 03-06 2015, Bellevue, WA, USA

ACM 978-1-4503-3969-8/15/11...$15.00
DOI: http://dx.doi.org/10.1145/2830834.2830838.

2. RELATED WORK

Many researchers have investigated methods to preserve geo-privacy. These methods can be categorized into two levels—coodinate and semantic. These respectively correspond to the explicit obfuscation of GPS coordinates and obfuscating at the semantic level, such as manipulating nearby location venues as introduced in Section 1. *K-anonymity* [19] is a typical method of coordinate obfuscation, where a user mixes his GPS coordinate with other $k-1$ nearby users' coordinates. These k coordinates are then all uploaded as one query. Alternatively, instead of providing an exact location, a user can cloak his coordinates and upload a geographical area instead [7]. Additionally, users can randomly perturb their coordinates by a limited distance and provide the moved position [7]. On the other hand, semantic obfuscation is implemented on a feature vector that has been extracted from absolute geo-coordinates as a step in the LBS processing pipeline. Semantic obfuscation by generalization is introduced in [4]. There, a specific venue (e.g. burger joint) may be replaced by a more generic venue category (e.g. restaurant) to mitigate the risk of an attacker reverse-engineering a user's absolute location. In our work, we contribute to the genre of semantic obfuscation approaches. This allows us to also consider LBS quality while preserving geo-privacy.

To quantify obfuscation performances, appropriate metrics are needed to correspond to the attacking scenario and LBS application context. Some basic metrics are proposed in [14]. Hoh and Gruseser [13] define *location privacy* as the expected error in distance between a user's true location and an attacker's estimate of that location. Duckham and Kulik [11] define *level of privacy* as the number of different geographic coordinates provided by the user in a single query. Beresford and Stajano [3] use entropy as the *privacy indicator*. Shokri et. al. put forward a *location-privacy meter* [17, 18] in consideration of the distance between attacker-estimated locations and real location, probability of attacker-estimated locations and the user's profile. In these metrics, *k-anonymity* obfuscation is appropriately quantified by *level of privacy*. In our work, we construct *k-anonymity* and *hit-rate* to appropriately quantify the level of geo-privacy preservation in our application context.

3. METHODOLOGY

Figure 1 demonstrates the assumed LBS working scenario when our obfuscation approach is utilized in the form of a trusted third-party, which has been successfully realized in practice [2, 8, 15]. In the figure, we demonstrate the positioning of the relevant parties: users, attacker, trusted third-party providing obfuscation, and LBS provider. The flow of information is as follows: A geographical region is first partitioned into unit cells via a spatial segmentation method. Once users request LBS, a trusted third-party receives the original geo-coordinates and converts them into semantic geo-features. Typically, these are the semantic categories of nearby venues extracted from the corresponding unit cells. Then, the semantic obfuscation approaches are applied and the obfuscated geo-feature vector is sent to the appropriate LBS provider. Our assumption is that the attacker has full access to the LBS provider and is able to intercept all the information transmitted from the trusted third-party to the LBS provider. We also assume that the geographic data is not linked with other third-party data, such as the credit card usage history of the corresponding user. Therefore, the geo-privacy will be truly protected given an effective geographic obfuscation approach. In addition, the third-party service is assumed to be running on the user's device for the sake of good communication quality. The following subsections detail each of the components in Figure 1 further.

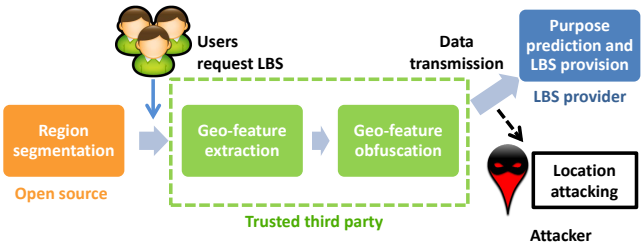

Figure 1: Location-based service (LBS) working scenario. The scenario includes a segmented region map, a trusted third party to extract original features and implement location obfuscation, an LBS provider for travel purpose prediction, and an attacker who attempts to locate the absolute geo-coordinates of the user.

3.1 Region Segmentation

To define an area from which the original geo-feature is extracted, we segment the region into unit cells. A typical method is to implement a grid-based segmentation as shown in [1]. In this work, Abdulazim et. al. developed an algorithm to adjust the grid size according to the venue density. However, restricted by the geometric shape, grid-based segmentation cannot capture natural spatial compositions as defined by population density, geographical terrain, or urban planning elements (e.g. roads, railways, etc.).

To tackle these problems, we employ a census-based segmentation method to realize a more natural space partitioning. Developed by the U.S. Census Bureau, census-based segmentation cells are the smallest units to collect, tabulate and present census and other geographic data[3]. Boundaries are defined by both physical features such as streets, streams, and railroad tracks as well as invisible boundaries such as town limits, property lines and imaginary extensions of streets and roads. An illustrative figure showing the segmentation of the Puget Sound area in the United States is shown in Figure 2.

3.2 Geo-feature Extraction

Based on the segmented region map, original geo-coordinates are assigned into unit cells that enclose them. Venues lying within the corresponding unit cells are extracted as semantic geo-features. We arrange the semantic geo-feature as a K-dimension vector $V = [v_1, ..., v_K]$, where K is the total number of venue categories and each dimension is a count of the k-th venue within the cell. As each trip has a starting and an end coordinate, original starting and end feature vectors are simultaneously constructed to characterize a trip. In our experiments, these non-obfuscated features also serve as the benchmark for comparison.

[3]http://www.psrc.org/data/gis/shapefiles

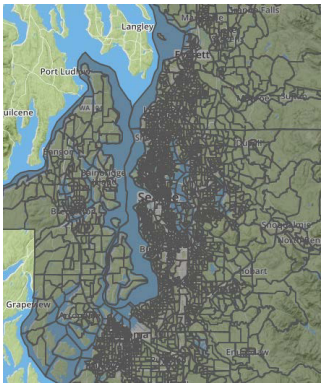

Figure 2: A part of segmented Puget Sound area. Cells are adjusted in size and shape based on geographic and demographic characteristics of the region.

3.3 Geo-feature Obfuscation

Obfuscation approaches are applied after the construction of benchmark feature vectors. Here, we propose three semantic-level obfuscation approaches.

3.3.1 Binarization Obfuscation

Benchmark feature vectors imply *which* and *how many* venues are included. Then, a simple obfuscation technique is to binarize the feature vector to eliminate the discriminability caused by venue counts. Consequently, we are left only with *which* venues exist. We choose the binarization threshold as 0, such that the dimension v_k will be 1 as long as at least one count of the corresponding venue exists in the cell. Figure 3 demonstrates an example of binarization obfuscation.

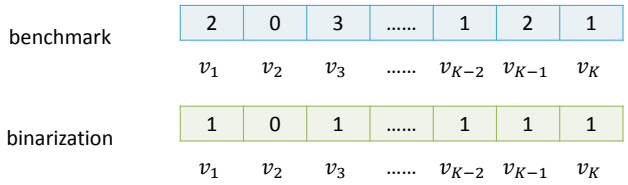

Figure 3: Blue vector is the benchmark and green vector is the result of binarization obfuscation. v_2 is not contained in the cell so its value is 0 while the other dimension values are 1.

3.3.2 Revealing-Venue-Switching Obfuscation (RVS)

Some venue categories are more revealing of an absolute geographical location than others. For instance, if there exists only 1 *Cambodian Restaurant* and it is in the user's feature vector, an attacker can easily locate the census cell where the user stays by geo-coding via an urban gazetteer. On the contrary, *Home* is relatively ubiquitous, therefore, barely reveals the user's location. To lower the risk of being exposed by revealing venues, we propose the *revealing-venue-switching obfuscation (RVS)*.

Given a region, we first calculate the degree of spread for each venue category according to the equation

$$s(K) = \sum_{i=1}^{K} \sum_{j=1}^{K} dist(k_i, k_j) \ for \ i \neq j \tag{1}$$

where $dist(k_i, k_j)$ is the geographic distance between two venues of category K in the region. The revealing rank is assigned inversely with the geographical spread: a venue category with the lowest degree of spread is ranked highest in revealing rank and vice versa.

During obfuscation, a revealing rank threshold is selected as a method parameter. A venue which ranks higher than this threshold will be replaced randomly by a less revealing venue with the same count. For example, 2 *Hockey Fields* may be switched to 2 *Parks*. Therefore, after RVS, total venue quantity is maintained and we call this *conservation of vector magnitude*. Due to random replacement, RVS does not produce unique output. Figure 4 illustrates an example explaining this method.

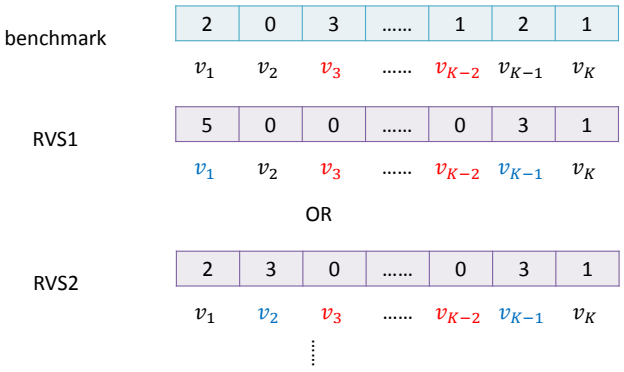

Figure 4: We assume that v_3 and v_{K-2} are revealing and all the others are unrevealing. We show 2 possible outputs RVS1 and RVS2. In RVS1, the revealing venues are switched to v_1 and v_{K-1}; while in RVS2, the revealing venues are switched to v_2 and v_{K-1}. The number of revealing venues are added to the corresponding unrevealing venues, which maintains the conservation of vector magnitude.

3.3.3 Binarized Revealing-Venue-Switching Obfuscation (bi_RVS)

Binarization eliminates the discriminability of venue count while RVS gets rid of revealing venues. As these two methods provide orthogonal obfuscation benefits, we combine them to incorporate the obfuscation effects of both. An example is illustrated in Figure 5.

3.4 Quality of Service from LBS

Understanding user intention is the essential ingredient in providing context-specific LBS. As a proxy assessment for LBS quality, we adopt the travel purpose prediction module in [20] to examine how our obfuscation methods degrade prediction accuracy. There, a multi-class classification problem is posed where supervised machine learning algorithms attempt to map spatial features to user intention for unseen trajectory instances. We utilize the same construction, aside from using a Random Forest classifier [6] in place of

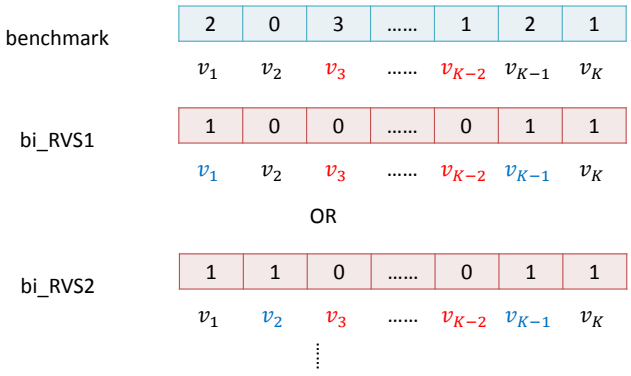

Figure 5: This example is based on Figure 4. We show 2 possible outputs and the binarization further obfuscates the feature vector by mapping it to 0-1 sequence.

the original Linear SVM classifier [12] due to its superior performance in this case.

3.5 Attacker Settings and Exclusion Rules

Typically, a LBS is requested upon arrival at the destination. Therefore, we assume the attacker's aim is to locate the trip end-point. In our purposes, a user is denoted as *attacked* as long as the attacker finds the cell the user is in. As the attack happens in real time during the data transmission (Figure 1), the attacker can obtain **the obfuscated spatial feature** and **upload timestamp**. We also assume that the attacker knows the **segmented region map**, **obfuscation method** and **corresponding parameters**.

Given the region map, the attacker will check each unit cell and target the user by reverse-engineering the implemented obfuscation approach. For each obfuscation approach, we define corresponding *exclusion rules* to filter out cells which have no possibility of containing the user. The following notations are used:

V^{obfs}–Obfuscated venue feature vector intercepted by attacker
V^{map}–Vector of venues from map unit cell
V^{bin}–Vector of venues from map unit cell after binarization
$V^{nr.map}$–Vector of non-revealing venues from map unit cell
$V^{nr.obfs}$–Vector of non-revealing venues from obfuscated venue features
$S^{nr.map}$–Set of non-revealing venues from map unit cell
$S^{nr.obfs}$–Set of non-revealing venues from obfuscated venue features
$S^{r.map}$–Set of revealing venues from map unit cell
$S_{diff} = S^{nr.obfs} - S^{nr.map}$–Relative complement of $S^{nr.map}$ with respect to $S^{nr.obfs}$

3.5.1 Binarization Exclusion Rule

- Exclusion Criteria:

$$\exists v_i^{bin} : v_i^{bin} \neq v_i^{obfs} \ for \ 1 \leq i \leq K$$

Exclusion rule for binarization is straightforward. After binarization, a unit cell can only remain if each of its dimension is the same as corresponding dimension in the intercepted vector. Otherwise it will be excluded.

3.5.2 RVS Exclusion Rule

- Exclusion Criteria 1:

$$\sum_{i=1}^{K} v_i^{map} \neq \sum_{i=1}^{K} v_i^{obfs}$$

This criteria reflects the *conservation of vector magnitude* as mentioned in section 3.3.2. If a cell feature vector meets this criteria, it will be excluded.

- Exclusion Criteria 2:

$$\exists v_i^{nr.map} : v_i^{nr.map} > v_i^{nr.obfs} \ for \ 1 \leq i \leq K$$

Given conservation of vector magnitude, the attacker can further compare venue feature values. If there *exists* non-revealing venues of the map cell larger than the corresponding venues in the obfuscated vector, it will be excluded.

3.5.3 bi_RVS Exclusion Rule

- Exclusion Criteria 1:

$$S^{nr.map} \not\subseteq S^{nr.obfs}$$

If before obfuscation, the map unit cell has already contained other non-revealing venues, which are not shown in obfuscated venue features, it cannot be obfuscated into the intercepted feature vector and should be excluded. If a cell is not filtered out by criteria 1, the attacker will further check it according to exclusion criteria 2.

- Exclusion Criteria 2:

$$card(S_{diff}) > card(S^{r.map})$$

$card(Set)$ is the count of distinct elements in a set. If a unit cell contains n different revealing venues, then its obfuscation will contain at most n more non-revealing venues. Therefore, if a cell meets criteria 2, it will be excluded.

After filtering according to exclusion rules, remaining cells form a candidate pool and they are regarded as potential user cells. To illustrate the exclusion rules more visually, a concrete example is provided in the Discussion section.

3.6 Location Attacking–Obfuscation Metrics

Based on the attacker candidate pool, we propose two obfuscation metrics *k-anonymity* and *hit-rate* to evaluate the effect of our obfuscation methods.

3.6.1 K-Anonymity

Originally, k-anonymity means the number of coordinates provided by the user to hide his/her location [19]. In our scenario, we utilize this term to represent the number of remaining cells in the attacker's candidate pool after applying exclusion criteria. The attacker will find it more difficult to locate the user if there are more potential user cells. Therefore, we pursue a large k-anonymity value.

3.6.2 Hit-rate

Previous works [13, 17, 18] adopt the theory that the user feels more threatened if the attacker-estimated locations are closer. Therefore, we introduce the metrics *hit-rate* @ Δ to

Purpose Label	Frequency	Examples
Go Home	34.64%	go back home or residential building
Work/Work-related	16.27%	go to meeting, delivery, etc.
Shopping	12.07%	go to grocery store, mall, pet store, etc.
Education	3.40%	go to elementary school, college, etc.
Professional Services	2.17%	go to the doctor, dentist, etc.
Personal Business	6.39%	go to bank, post office, etc
Drop off/Pick up Someone	6.08%	pick up children after school, etc.
Exercise	5.50%	go to gym,bike-ride, etc.
Eat Out	5.89%	go to restaurant, get take-out, etc.
Socializing	3.17%	visit with friends, co-workers, etc.
Recreation	2.05%	go to movies, sporting centers, etc.
Community Activity	1.18%	go to church, volunteer activity, etc.
Transfer to Another Mode of Transportation	0.19%	change from ferry to bus, etc.
Other	0.93%	activities not in items above

Table 1: 14 Travel Purposes of PSRC Survey and Their Occurrence Frequencies

calculate the proportion of attacker cells within Δ meters of the user's actual location. It is formulated as:

$$HR(\Delta) = \frac{N_N(\Delta)}{N_T(\Delta)} \quad (2)$$

where N_N is the count of cells within Δ meters of the user and N_T is the total count of cells remaining in the attacker's candidate pool. Therefore, we pursue a low hit-rate value.

4. EXPERIMENTAL RESULTS

According to the scenario detailed in the last section, we evaluate the performance of our obfuscation approaches regarding the geo-privacy preservation with consideration of quality-of-service of LBS. The LBS prediction accuracy is the weighted test-fold classification over 10-fold cross-validation. As in [20], the folds are generated carefully so that they are not shared by a single user, in case of repeated trips, which may artificially boost the accuracy. As we implement random switching in RVS, the results pertaining to RVS are the average over 20 iterations. We begin this section by first introducing the experimental dataset.

4.1 Dataset

We use the Puget Sound Regional Travel Study conducted by the Puget Sound Regional Council (PSRC) of Washington State, USA in the year of 2014. This survey contains 47881 trips in a 24-hour weekday period from 6094 representative households including around 10K people [10]. The respondents come from different occupations with age ranging from *under 5 years old* to *85 or older*.

This survey investigate the travel purpose for each trip as explicitly indicated by participants. Originally, 16 purposes are included and we slightly reduce the count to 14 by combining *go to workplace* and *go to other work-related place* as *go to work/work-related place*, and by combining *go grocery shopping* and *go to other shopping* as *go shopping*. Detailed examples and occurrence frequencies of travel purposes are summarized in Table 1.

For spatial feature extraction, we collect the crowd-generated venue data from the social media platform Foursquare. With the help of Venues API service from the platform[2], we populate the Puget Sound region with 111725 venues which are categorized into 592 semantic categories such as *College Academic Building*, *American Restaurant*, *Church*, etc.[3] Since 22 venue categories appear only once, we only have 570 revealing rank levels. Table 2 shows typical venue categories in different rank intervals.

Revealing Ranks	Typical Venue Categories
1-100	Monastery, Palace, Tibetan Restaurant, Volcano, Hindu Temple, etc.
101-300	Whiskey Bar, Shrine, German Restaurant, Prison, TV Station, etc.
301-500	Zoo, Supermarket, Burger Joint, Credit Union, Baseball Field, etc.
501-570	Bus Stop, Church, Elementary School, Office, Home(private), etc.

Table 2: High-rank venues are considerably sparse in the region. For example, there is only one Tibetan restaurant in Puget Sound area. Low-rank venues are relatively ubiquitous such as office and home.

In terms of the geographical region covered by the PSRC dataset, 28059 segmented census cells contain Foursquare venues. From these, 8523 cells contain end coordinates of 35725 trips. Another 12156 trips are located in cells without venues, therefore, we do not consider them in our experiments as feature extraction for these are not possible.

4.2 Benchmark Performance

As mentioned in Section 3.3, actual semantic geo-features are used in benchmark experiments. Using these features, the average test-fold prediction accuracy is 60.35%. Considering that we only use semantic geo-feature to classify 14 purposes which are distributed unevenly, this accuracy is satisfying compared with [1], where 6 classes are classified with a 67% accuracy. However, the privacy preservation of non-obfuscated geo-features is quite poor as shown in Figure 6. In hit-rate calculation, we set the hit distance threshold (Δ) as $50m$, $200m$, $500m$, $1000m$ and $2000m$. As the resultant average hit-rates increase within 2%, we only report the result at $\Delta = 500m$ in this paper. At this threshold, the average k-anonymity is 12.52 while the average hit-rate is 88.45%. To obtain a more representative view of performance for all trip instances, we plot the cumulative distribution function (CDF) of k-anonymities and hit-rates of 35725 trips in Figure 6.

[2]https://developer.foursquare.com/overview/venues.html
[3]http://developer.foursquare.com/categorytree

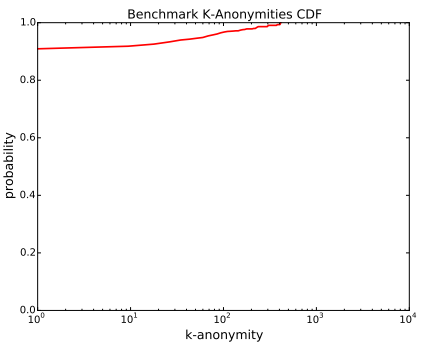

(a) Benchmark K-anonymities CDF

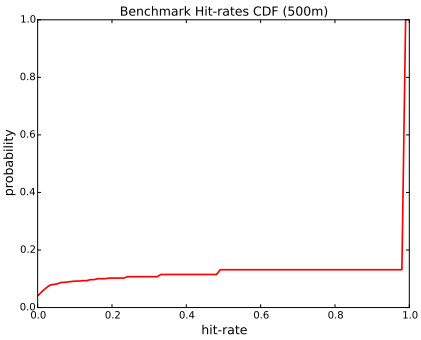

(b) Benchmark hit-rates CDF

Figure 6: Culmulative distribution functions (CDF) of benchmark k-anonymity and hit-rate. About 90% of trips have a k-anonymity of 1 and hit-rate of 100%, which means that the user location is totally exposed.

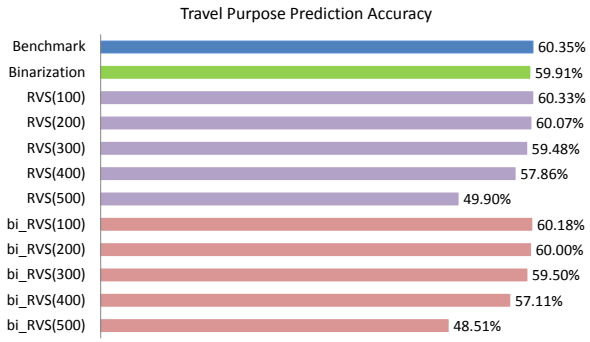

Figure 7: Purpose prediction accuracy of benchmark, binarization, RVS and bi_RVS. Numbers in brackets are revealing rank thresholds. Compared with benchmark, except for RVS and bi_RVS at revealing rank threshold 500, the biggest accuracy drop is only 3.24%.

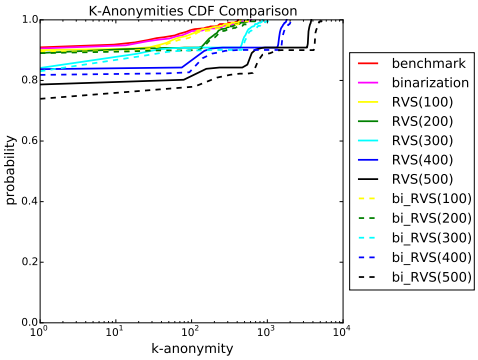

(a) K-anonymities CDF Comparison

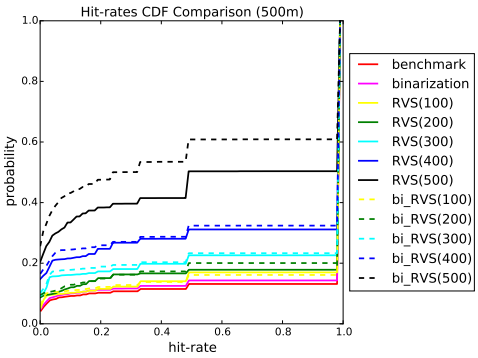

(b) Hit-rates CDF Comparison

Figure 8: Compared with benchmark, binarization barely improves geo-privacy preservation; RVS and bi_RVS improve geo-privacy preservation gradually as revealing rank threshold increases. On the whole, bi_RVS at high rank thresholds works significantly as x-axis of k-anonymity is in log scale and the hit-rate curve rises by approximately 40%.

In Figure 6, we demonstrate that around 90% of trips have k-anonymity of 1 and near 100% hit-rate. The poor geo-privacy preservation performance substantiates the necessity to implement obfuscation on the semantic geo-features.

4.3 Obfuscation Performance

We first compare the prediction accuracies of the obfuscation methods in comparison with the benchmark in Figure 7. Here, we demonstrate relatively small drops in accuracy after obfuscation methods are applied. Except for RVS and bi_RVS at revealing rank threshold 500, the largest accuracy drop is only 3.24% when implementing bi_RVS(400).

We then compare the geo-privacy preservation performances as shown in Figure 8. The best privacy preservation is achieved when applying bi_RVS, especially at threshold 400 in consideration of the travel purpose prediction accuracy. Compared with benchmark where around 90% of 35725 trips have a hit-rate higher than 95%, only 50% of the trips have such a high hit-rate when applying bi_RVS(400). On average, at this rank threshold and compared to benchmark, k-anonymity is 190.51 which increases by 15.22 times, and hit-rate is 70.90% which decreases by 17.55%.

However, the geo-privacy preservation improvement is re-

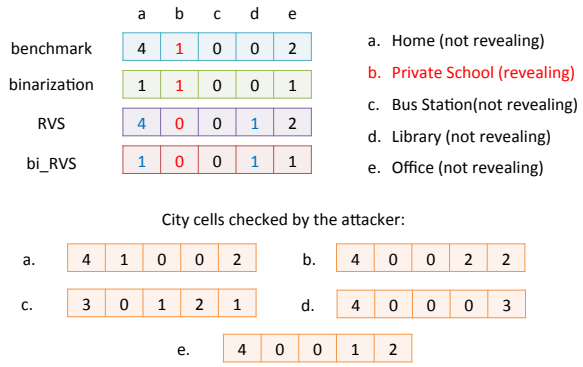

Figure 9: Venue b is assumed to be a revealing venue according to the rank threshold. 3 obfuscated feature vectors are listed behind the benchmark. 5 feature vectors from different unit cells are checked by the attacker to find the real user cell.

stricted by the feature vector specificity in the placement of venues. 32277 out of 35725 trips end in cells which have a unique feature vector and manipulation of these vectors hardly leads to improvements. The obfuscation can be more effective if the unit cells are generated in a specificity reduction manner.

5. DISCUSSION

In the last section, we show that RVS and bi_RVS make significant gains in geo-privacy preservation while maintaining a high accuracy in travel purpose prediction. We provide some intuitive discussions as to how our methods are capable of achieving the demonstrated results.

5.1 Geo-privacy Preservation

Section 3.5 details the exclusion rules for each obfuscation approach. In this part, we illustrate a concrete example in Figure 9. In this example, we assume that 5 venues exist in the region for simplicity.

- *Benchmark*

As the attacker knows that the benchmark is uploaded, he will match only map cells with the same feature vector. In Figure 9, only *cell a* meets the requirement. The feature specificity mentioned in the last section explains the poor geo-privacy preservation of benchmark.

- *Binarization*

Similar to benchmark, binarization hardly improves the geo-privacy preservation level due to the feature specificity. All cells except for *cell a* will be filtered out from the attacker candidate pool.

- *RVS*

First, *cell b* will be excluded because it meets exclusion criteria 1. Then, according to criteria 2, the attacker can filter out *cell c* and *cell d*. *Cell a* and *cell e* will remain as potential locations to the attacker. RVS improves the geo-privacy preservation level gradually with the increase of revealing rank thresholds. However, specificity of venue counts restricts this approach.

- *Binarized RVS*

According to exclusion criteria 1, *cell c* will be filtered out. Then, *cell d* will be excluded due to criteria 2. Finally, *cell a*, *b* and *e* remain in the candidate pool. It can be seen that binarization eliminates the influence from venue count, which leads to better geo-privacy preservation compared with pure RVS. However, exclusion criteria 2 is relatively strict and still aids the attacker in filtering out potential user cells.

5.2 High LBS Quality after RVS Obfuscation: Revealing Is Not Informative

RVS and bi_RVS maintain a relatively high prediction accuracy even after switching a fair number of venue categories. Even at rank threshold 400, where around 70% of the revealing venue categories are replaced, the prediction accuracy is still high at 57.86% (without binarization) and 57.11% (with binarization). In Figure 10, we plot the revealing rank on the x-axis while the normalized feature importance (via Gini entropy importance [5]) on the y-axis. From the figure, we see a salient pattern: venues *most informative* of the purpose are also some of the *least revealing* ones. Therefore, although RVS and bi_RVS switch out the revealing venues, LBS quality is not significantly influenced as those are likely ignored by the classifier as uninformative venues.

6. CONCLUSION AND FUTURE WORK

In this work, we demonstrate the usefulness of pure semantic geo-feature in travel purpose prediction. Our work explores location obfuscation methods and presents a *Binarized Revealing-Venue-Switching Obfuscation (bi_RVS)* that achieves good geo-privacy preservation while maintains LBS quality. Compared with the benchmark, the travel purpose prediction accuracy is impaired slightly by 3.24% but the obfuscation level improves by 17.55% in hit-rate and 15.22 times in k-anonymity.

Future work includes adding in temporal and demographic information, which would help increase the travel purpose prediction accuracy, thus, improve the LBS quality. Correspondingly, we could extend our methods to provide more robust and comprehensive obfuscation approach by considering these newly-added elements. Also, due to feature specificity of some urban regions, designing obfuscation that prioritizes popular destination cells would be helpful. For example, segmenting geo-regions in a feature-specificity-free style would help decrease specificity of semantic features.

7. REFERENCES

[1] T. Abdulazim, H. Abdelgawad, K. M. Nurul Habib, and B. Abdulhai. A framework to automate travel activity inference using land-use data: The case of foursquare in the greater toronto and hamilton area. In *Transportation Research Board 94th Annual Meeting*, number 15-5850, 2015.

[2] B. Bamba, L. Liu, P. Pesti, and T. Wang. Supporting anonymous location queries in mobile environments with privacygrid. In *Proceedings of the 17th international conference on World Wide Web*, pages 237–246. ACM, 2008.

[3] A. R. Beresford and F. Stajano. Location privacy in pervasive computing. *IEEE Pervasive computing*, 2(1):46–55, 2003.

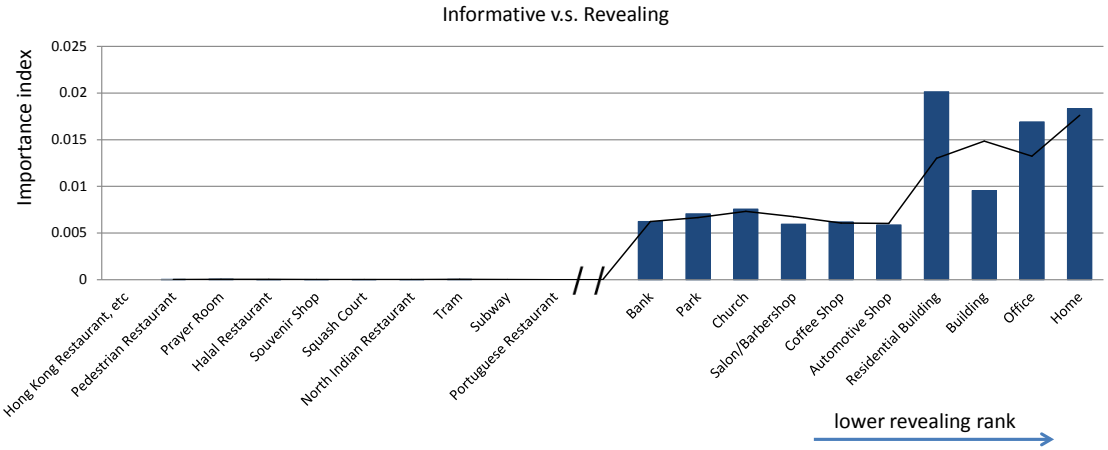

Figure 10: Venue features are arrayed in the ascending order regarding the revealing rank on the x-axis, i.e. the most revealing venue is on the left while the least revealing venue is on the right. The y-axis represents the venue importance index for the classifier. This graph demonstrates that the classification-important venues are actually not revealing in the region and vice versa.

[4] I. Bilogrevic, K. Huguenin, S. Mihaila, R. Shokri, and J.-P. Hubaux. Predicting users" motivations behind location check-ins and utility implications of privacy protection mechanisms. In *22nd Network and Distributed System Security Symposium (NDSS" 15)*, number EPFL-CONF-202202, 2015.

[5] L. Breiman. Technical note: Some properties of splitting criteria. *Machine Learning*, 24(1):41–47, 1996.

[6] L. Breiman. Random forests. *Machine learning*, 45(1):5–32, 2001.

[7] A. Brush, J. Krumm, and J. Scott. Exploring end user preferences for location obfuscation, location-based services, and the value of location. In *Proceedings of the 12th ACM international conference on Ubiquitous computing*, pages 95–104. ACM, 2010.

[8] C.-Y. Chow and M. F. Mokbel. Privacy in location-based services: a system architecture perspective. *Sigspatial Special*, 1(2):23–27, 2009.

[9] C.-Y. Chow, M. F. Mokbel, and X. Liu. A peer-to-peer spatial cloaking algorithm for anonymous location-based service. In *Proceedings of the 14th annual ACM international symposium on Advances in geographic information systems*, pages 171–178. ACM, 2006.

[10] P. S. R. Council. 2014 household activity survey: General user's guide. In *http://www.psrc.org/assests/12061/2014-Household-Survey-Dataset-Guide.pdf*, 2014.

[11] M. Duckham and L. Kulik. Simulation of obfuscation and negotiation for location privacy. In *Spatial Information Theory*, pages 31–48. Springer, 2005.

[12] R.-E. Fan, K.-W. Chang, C.-J. Hsieh, X.-R. Wang, and C.-J. Lin. LIBLINEAR: A library for large linear classification. *Journal of Machine Learning Research*, 9:1871–1874, 2008.

[13] B. Hoh and M. Gruteser. Protecting location privacy through path confusion. In *Security and Privacy for Emerging Areas in Communications Networks, 2005. SecureComm 2005. First International Conference on*, pages 194–205. IEEE, 2005.

[14] J. Krumm. A survey of computational location privacy. *Personal and Ubiquitous Computing*, 13(6):391–399, 2009.

[15] M. F. Mokbel. Towards privacy-aware location-based database servers. In *Data Engineering Workshops, 2006. Proceedings. 22nd International Conference on*, pages 93–93. IEEE, 2006.

[16] J. P. Munson and V. K. Gupta. Location-based notification as a general-purpose service. In *Proceedings of the 2nd international workshop on Mobile commerce*, pages 40–44. ACM, 2002.

[17] R. Shokri, G. Theodorakopoulos, J.-Y. Le Boudec, and J.-P. Hubaux. Quantifying location privacy. In *Security and Privacy (SP), 2011 IEEE Symposium on*, pages 247–262. IEEE, 2011.

[18] R. Shokri, G. Theodorakopoulos, C. Troncoso, J.-P. Hubaux, and J.-Y. Le Boudec. Protecting location privacy: optimal strategy against localization attacks. In *Proceedings of the 2012 ACM conference on Computer and communications security*, pages 617–627. ACM, 2012.

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

[20] Z. Zhu, U. Blanke, and G. Tröster. Inferring travel purpose from crowd-augmented human mobility data. In *Proceedings of the First International Conference on IoT in Urban Space*, pages 44–49. ICST (Institute for Computer Sciences, Social-Informatics and Telecommunications Engineering), 2014.