Sensitivity Analysis of Personal Location Disclosure

John Krumm Microsoft Research Microsoft Corporation Redmond, WA USA jckrumm@microsoft.com

Abstract—Regular people give away their location data without much knowledge about what can be inferred from it. This paper presents a sensitivity analysis of location disclosure, showing what can be inferred from even just a few place visits. The aim is to highlight how modern inference algorithms can take small location disclosures and create detailed inferences about where else a person is likely to go and what sorts of places they tend to visit. Based on an analysis of over 100,000 people, we show how disclosing just one location point can be used to predict other visits with an AUC of 0.88. We develop another algorithm that shows the types of places a person has an unusually high propensity to visit, and we use this result as part of an economic analysis of delivering targeted advertising. This research serves to raise awareness about what can be inferred from even very small location disclosures, which can in turn inform regular people about their true privacy risks.

Index Terms—personal location, location privacy, inference, advertising

I. INTRODUCTION

Individuals are largely unaware of how their location data is used to make inferences about them. A recent survey showed that less than one third of consumers understand how retailers use their personal data after it is disclosed [1]. Even understanding who has copies of the data is difficult to know, as data brokers buy, sell, and trade personal data [2]. Some of these brokers specialize in location data, gathering the whereabouts of over 200 million Americans from apps on their phones [3]. Intellectual property concerns are a disincentive for companies to reveal the details of how and what they can infer from personal data. A legal analysis of Europe's GDPR privacy law suggests it does not necessarily guarantee an individual can successfully demand a detailed explanation of an inference based on their personal data, much less an explanation of which inferences may be made and how [4].

The research community can help in this regard by showing what can be inferred from personal data, including location data. This paper is intended to demonstrate the types of inferences that are possible from even a small location disclosure from an individual, such as just a few discrete locations. In particular, from a small disclosure, we demonstrate a pipeline of inferences that reveal (1) the individual's likely other location visits, (2) their propensity to visit locations with certain categories of businesses, and (3) the economic reasoning that a potential advertiser would use to determine whether or not to deliver an ad to the individual based on the the inferences in (2). This paper presents a plausible sequence



Fig. 1. The blue triangles show the discrete visit regions covering the extent of our experimental location data in the suburbs east of Seattle, WA. Thicker outlines indicate more popular visits averaged over all our location data.

of inferential steps that a disclosure recipient could employ to take advantage of the disclosure.

We refer to this investigation as a sensitivity analysis, because it considers the case of disclosing a small amount of location data compared to disclosing none. It is essentially approximating the derivative of inference accuracy vs. the size of the disclosure when the existing disclosure is small or zero. This is the situation a user faces when an app first asks for permission to access their location data. Our analysis shows what inferences could result from granting this permission after only a few location disclosures, i.e. the sensitivity of disclosures at the point of little or no previous disclosures.

We aim to identify and communicate specific inference implications. Revealing these implications are not mandated by law [4] nor likely to be divulged voluntarily by enterprises making the inferences. This paper develops and describes principled approaches to making these inferences and shows how well they work on data from about 100,000 people who made about 2.8 million discrete location visits over five months. To the best of our knowledge, this question of inference sensitivity has not been previously addressed in the research literature, because previous work does not address small location disclosures of individuals' everyday lives.

II. RELATED WORK

With a goal of understanding how small location disclosures can lead to significant inferences, this work is related to multiple threads of previous research: location privacy, location prediction, and data completion.

A. Location Privacy

The goal of this paper is to understand and report what can be inferred about a person with only a small location disclosure. The privacy paradox says that while people say they are concerned about privacy, including location privacy, they do not actively protect their personal data and often knowingly give it away [5]–[7]. For location privacy in particular, a survey by Zafeiropoulou et al. showed that, out of 150 respondents, 79% professed concern about location privacy, but they survey did not find a strong correlation between the respondent's privacy concerns and their actual behavior [8].

One way to induce privacy-enhancing behavior could be to reveal what sort of inferences can be made about a person from their location data. Almuhimedi investigated this in his PhD thesis [9]. His survey asked 861 respondents whether or not they would adjust their mobile phone's privacy settings in response to privacy "nudges" that highlighted increasingly alarming consequences of certain apps accessing the phone's location data. The results showed that users were more likely to adjust their privacy settings when they were told about potential inferences that could be made about them from their location data. For instance, one statement said, "These apps have access your location in the past week. With this information, apps can infer additional details about you, such as the address where you live, and use it to predict your income." This induced statistically significant more willingness to adjust privacy settings over the baseline statement of "These apps have accessed your location in the past week." This implies that users may be more willing to address location privacy if they understand what can be inferred about them from their location data. The present paper aims to clarify the inferential possibilities from even a small location disclosure.

B. Location Prediction

The research literature in predicting a person's location is long, with some of the earliest work by Marmasse & Schmandt [10] and Ashbrook & Starner [11] using a Markov model among other techniques. Song et al. explore the ultimate limits of human location prediction [12]. We are particularly interested in understanding prediction sensitivity based on only a small set of disclosed locations, which narrows the field of related work.

While this paper looks at this problem from a novel sensitivity point of view, there is related work in predicting the movements of tourists who naturally have little location history of the places they visit. Chen et al. focus on using tourists' four previous stops to predict the fifth using a variety of methods, finding that an LSTM worked best [13], with an accuracy of 94.8%. The test data was from the country of Andora, and the prediction was made among 13 different regions using not only previous locations but other features such as the nationality of the tourist, weekday/weekend, and points of interest (POI) in each region. Using 432 discrete cell locations, Zheng et al. deployed a variable-order Markov model to predict tourists' next destinations at Beijing's Summer Palace, achieving an accuracy of 69.8% [14]. In tourist prediction, the work most closely related to ours comes from Muntean et al. [15]. They study the next location prediction problem using a rich set of features describing each tourist's previous visits. Part of their study examines the accuracy of their ranking model as a function of how many previous visits are included in the tourist's history for prediction, ranging from one to seven. Interestingly, accuracy seems to flatten after the history length reaches about three previous visits.

Our work differs from these efforts in tourist prediction in that we infer *all* visits, past and future, not just the next, from between one and five previously disclosed locations. Also, our output space is relatively large, at 3218 possible location cells (Figure 1). This significantly exceeds the previous work of 13, [14](432), and [15](888). These differences are due to the different nature of our problem, where our goal is to infer *all* visited locations of a person in their normal life, which is not limited to popular tourist destinations. We are especially interested in the accuracy of these inferences based on only a small amount of location history in an effort to understand the sensitivity of even small location disclosures for inference-based privacy attacks.

C. Data Completion

Our goal is to understand the privacy risk of a few location disclosures, i.e. visits to a few cells in Figure 1. One of the ways we measure risk is by inferring other visits from a small amount of disclosures. This is related to the general problem of data completion. One of the most general forms of this problem is matrix completion or the more general tensor completion [16]. In this scheme, the data is normally represented as a multidimensional array of numbers, either a matrix or a tensor, with some elements missing. The general approach is to find a complete, low-ranked array such that its elements match the known elements of the original array. These problems are characterized by their "missing ratio," which is the fraction of missing data in the array. For our problem, we look at disclosures with only one to five locations out of all the cells in Figure 1, leading to a missing ratio of around 99.8%, normally considered extremely high for traditional data completion algorithms.

Matrix completion plays a prominent role in some recommender systems [17], such as entries for the Netflix Prize [18]. In a recommender system, the problem is to compute recommendations for new items based on a person's previous ratings. Our problem is similar if the items are visits to locations or POI. It is notable, however, that visiting a certain place does not necessarily imply a positive recommendation [19]. One of the algorithms we test, k nearest neighbors, is commonly used for recommender systems.

Another related research effort is image inpainting or image completion [20]. Starting from an image with missing or corrupted pixels, the goal is to fill the pixels to make a reasonablelooking, complete image. Current solutions are dominated by deep neural net models that can learn from huge collections of images that have been artificially corrupted in some way. Our best algorithm for location inference, in Section IV-E, uses a neural net autoencoder, which is a common architecture for image inpainting.

III. LOCATION DATA

Our experimental data comes from Safegraph¹, which is a company that aggregates location data taken from mobile phones. This is an especially appropriate source of data for our work, because it represents the type of location data that is easily available to advertisers and other enterprises that want to understand where people go. Each location record consists of a UTC timestamp, latitude/longitude pair, precision, and a device identifier for the phone. We took data for five months staring in December of 2020, limited to the eastern suburbs of Seattle, WA shown in Figure 1. The total number of points per person in a typical month (April 2021) had a mean of 94.6, a median of 2, and a 90th percentile of 169.

The inferences we describe in the remainder of the paper pertain to visits to discrete locations on the map. We discretized space with the hierarchical triangular mesh (HTM) [21]. We used HTM level 15 whose triangles have an area of about 0.06 km^2 and a side length of about 0.5 km, shown in Figure 1. We chose this size because it represents a reasonably accessible area that a person may visit near a measured location point.

Ultimately we want to infer which other triangles a person may visit based on recorded visits to a small number of the triangles. We compute a triangle visit in a simple way. Each person's data is divided into 30-minute segments starting at midnight on the first day of their data. We compute the median latitude/longitude of the data in each 30-minute segment and declare a visit to the median point's enclosing triangle. We chose 30 minutes to represent a reasonable visit time. This basic approach to visit detection was necessitated by the wide variety of trajectory types: some people's data is sampled very sparsely, while others show occasional dense trajectories along apparent driving paths. The simplicity of this technique may generate false positive visits while a person is in transit, but the 30-minute interval avoids a dense sequence of inferred visits along a moving trajectory. In addition, these in-transit visits would be shared with many other visitors which are implicitly ignored when we look for unusual behavior in our POI analysis in Section V. Of the 3278 triangles in Figure 1, only 60 were unvisited. We disregard these 60 for the remainder of the paper and concentrate only on the $N_v = 3218$ that were visited.

For a person *i*, all their visits are represented as a set of n_i triangles $V_i = \{v_{i1}, v_{i2}, ..., v_{ij}, ..., v_{in_i}\}$. As a formal set, V_i contains no repeated elements and no consistent ordering, so we do not look at the frequency of visits nor their time stamps for simplicity. For meaningful inferences, we kept only those persons with $n_i = |V_i| \ge 5$, which reduced the number of persons from 653,203 to 101,507. Over the five month period, the total number of visits among the remaining set of persons was 2,796,346, and the median number of visits per person was 12.

Our goal in the next section is to estimate the full set V_i from a small subset of V_i .

IV. SENSITIVITY ANALYSIS OF LOCATION INFERENCE

We know that location is considered by many people to be a sensitive piece of personal data [21]. In this section we investigate how well we can infer *all* the places someone visits based on only a small subset of their visits. More specifically, for person *i*, we want to predict their set of actual visits V_i from a subset of their visits $V_i^{(s)}$, where the size of the subset is small, i.e. $|V_i^{(s)}| \in \{1, 2, 3, 4, 5\}$. Note that *s* is the size of the subset, i.e. $s = |V_i^{(s)}|$. This inference is designed to understand how much can be learned about where a person goes from only a small amount of their actual location data.

This section describes and evaluates five methods for this problem. Each method produces a visit probability for each triangular region given the visits already observed. This is $P(v_{ij} = 1|V_i^{(s)})$. The subscript *i* indicates the particular person. The subscript *j* indexes over all the possible visit locations, so $j \in \{1...N_v\}$, where N_v is the number of possible visit triangles. $N_v = 3218$ in our case. The binary visit variable v_{ij} has a value of one if person *i* visited location *j*, and zero otherwise. Naturally $P(v_{ij} = 0|V_i^{(s)}) = 1 - P(v_{ij} = 1|V_i^{(s)})$.

Our assumption is that an inference-maker would have access to past location histories of a substantial number of people to use for training their inference method. Unless otherwise noted, each method we described is trained on a random 90% of the people in our data and tested on the remaining 10%. We generated a separate set of training and test instances for each subset size $s = |V_i^{(s)}| \in \{1, 2, 3, 4, 5\}$. Recall that the subset is the set of visited locations that the inference can see in its attempt to infer *all* the person's visits. For each test person and each subset size s, we generated one random subset of s visited locations from their actual visited locations. Thus for each subset size, the number of test persons.

We will explain how we generated the training data in Section IV-E, which is the first method to use training data in the traditional way.

¹https://www.safegraph.com/

Each of the methods we explain is a binary classifier for whether or not the person has or will visit v_{ij} given that the classifier has seen visits to $V_i^{(s)}$. The classification decisions \hat{v}_{ij} come in the usual way from thresholding the inference probability $P(v_{ij} = 1|V_i^{(s)})$ by a threshold $0 \le T \le 1$, i.e.

$$\hat{v}_{ij} = \begin{cases} 0 & \text{if } P(v_{ij} = 1 | V_i^{(s)}) < T \\ 1 & \text{if } P(v_{ij} = 1 | V_i^{(s)}) \ge T \end{cases}$$
(1)

As a conventional classification problem, we evaluate the results with the area under the curve (AUC) of the receiver operating characteristic (ROC) curve. AUC is appropriate for this problem because it considers rates of false positives and true positives, rather than their absolute numbers, imposing a balanced view of the two. In our particular problem, the median number of ground truth visits was 12, giving a ground truth median positive rate of only $\frac{12}{N_v} = 0.00373$. The ground truth data is heavily biased toward negatives, because people do not generally visit many places.

A. Copy Input

This is a simple method that copies the input set of visits to the inferred set of output visits. It says that the person will visit only the places she has been observed to visit in the past. Mathematically

$$\hat{P}(v_{ij} = 1 | V_i^{(s)}) = \begin{cases} 0 & \text{if } v_{ij} \notin V_i^{(s)} \\ 1 & \text{if } v_{ij} \in V_i^{(s)} \end{cases}$$
(2)

In terms of performance, all the positive inferences from this method are true positives, because it only makes a positive prediction for visits that have actually been observed. This means its sole operating point on the ROC curve is on the vertical axis, with false positive rate of zero, as shown in Figure 2. In order to compute an AUC, we artificially extended each curve from its sole operating point to the upper right of the ROC curve. This upper right point could represent an approach that predicts *all* the triangles would be visited.

Table I shows the AUC over the test cases as a function of the size s of the observed subset $V_i^{(s)}$. As expected the AUC increases with more location disclosure, but the low AUC values show that this is not an effective method.

B. Shared Prior

This method infers $P(v_{ij} = 1|V_i^{(s)})$ from a shared prior over all the persons in the training set. We used the training data to estimate the prior probability of visiting location v_j , computed as the proportion of test users who visited v_j , called $P(v_j = 1)$. Thus this method says simply $P(v_{ij} = 1|V_i^{(s)}) =$ $P(v_j = 1)$. The prior $P(v_j = 1)$ is approximately the same as the visit proportions shown for the whole dataset in Figure 1.

The approach is based on the assumption that every person shares the same set of visit probabilities. While this is clearly not true, the method still performs fairly well with an AUC of 0.843. In this case, the inference is completely independent of the observed visits, so the AUC does not depend on the size

ROC Curves for Copy Input Method



Fig. 2. These are the ROC curves for the "Copy Input" method. It only gives true positives, so we artificially extended the curves to the upper right corner to compute the AUC values in Table I.

s of the observed subset $V_i^{(s)}$. From Table I, it is clear that the shared prior method significantly outperforms the "Copy Input" method in terms of AUC.

C. Joint Probability Distribution

From the training data we can estimate a joint probability distribution over all the N_v possible visits, $P(v_1, v_2, ..., v_{N_v})$ giving the approximate probability of any set of visits. This discrete distribution has N_v dimensions, and each dimension is two units long, representing the either a visit or not. From the disclosed visits we compute a conditional distribution of visits $P(v_1, v_2, ..., v_{N_v} | V_i^{(s)})$, which gives a new joint distribution. Technically the conditional distribution does not include the disclosed locations, so we write it as $P(\{v_1, v_2, ..., v_{N_v}\} \setminus$ $V_i^{(s)}|V_i^{(s)}$ to be precise, where $\{v_1, v_2, ..., v_{N_v}\} \setminus V_i^{(s)}$ does not include the disclosed locations. From the conditional distribution, we read out the "visit" and "not visit" probabilities for each v_i , normalize the pair so they add to one, and keep the normalized "visit" probability as $P(v_{ij} = 1)$. The probability of visiting the disclosed locations in $V_i^{(s)}$ is one. From these visit probabilities, we can compute the ROC and AUC as above, giving the results in Table I. Here we see that the joint probability method outperforms the first two methods ("Copy Input" and "Shared Prior") for small disclosures of s = 1 and s = 2, but underperforms "Shared Prior" for larger disclosures. In fact, its performance drops with larger values of s. This is counterintuitive, because we would expect more location disclosure would lead to better inferences. This underperformance is likely because the conditional PDF is based on fewer and fewer training samples as s increases. Creating the conditional PDF is essentially a process of finding those training examples whose set of visits is a superset of disclosed visit. As there are more disclosed visits, the number of qualifying training examples decreases, as shown in Figure 3. When s = 1 there are a median of 2632 relevant training samples, but this drops to 4 when s = 5. Essentially, the joint distribution has too many dimensions to estimate accurately with our data. We fix this with the next method.

TABLE I AUC VALUES FOR LOCATION INFERENCE METHODS

Dis- closure Size	Сору	Prior	Joint	kNN	MLP
1	0.518	0.843	0.880	0.879	0.869
2	0.536	0.843	0.879	0.892	0.890
3	0.554	0.843	0.826	0.900	0.903
4	0.572	0.843	0.809	0.909	0.912
5	0.590	0.843	0.766	0.913	0.920

Median Number of Matching Samples for Marginal Joint PDF



Fig. 3. For a given number of disclosed locations, this is the median number of samples in the training data that matched for the computation of the marginal joint PDF.

D. K Nearest Neighbors

The k nearest neighbor (kNN) algorithm finds training instances that are similar to the set of test visits $V_i^{(s)}$. We will represent a training vector as v. It is a binary vector with one element for each of the possible N_v visit locations. The corresponding set of positive visits in the training vector is V. The distance between the test visits and training visits is the number of positive visits that they share, i.e. $|V_i^{(s)} \cap V|$. kNN finds the k nearest training vectors and averages them to estimate the visit probabilities $P(v_{ij} = 1)$.

We tested different values of k and found k = 400 worked well, giving the AUC results in Table I. Except for nearly matching the performance of the joint PDF for s = 1, kNN outperformed all the previous methods in this paper. Its AUC also grows intuitively with s, which matches our expectations that more location disclosure leads to better inferences about other visits.

The kNN and joint PDF methods are similar in that they both select and average some subset of the training vectors. The joint PDF method selects only those training vectors that contain a superset of the test set V_i^s , which leads to the paucity of relevant training vectors described above. kNN instead averages a preset number of nearly matching training vectors, which in this case leads to better inferences.

E. Multilayer Perceptron

Our final location inference algorithm is a multilayer perceptron (MLP). This is a "vanilla" neural network whose architecture is shown in Figure 4. The input layer has N_v nodes, one for each visit location. For an input disclosure, all the input nodes are zero except for the nodes corresponding to the visits in $V_i^{(s)}$, which are one. The output layer has the same number of nodes as the input layer, also corresponding to the visit locations. In the training phase, these outputs are binary, representing the actual visits.

The neural net architecture is a simple autoencoder with a bottleneck of 10 nodes as the middle layer, shown in Figure 4. A traditional neural network autoencoder maps an input vector to itself, reducing the dimensionality of the input to a relatively small number of nodes in the hidden layer(s). In our case, the input is an "s-hot" binary vector of the N_v possible visit locations in $V_i^{(s)}$, with the visited locations represented by a one and the remaining unknown-visit locations represented by a zero. In training, the output is a binary vector of the same size representing all the visited locations of person i. The latent space of the hidden layer is designed to learn a compact representation of typical visit patterns that is detailed enough to be accurate but general enough to ignore outliers. The activation functions exiting the input and hidden layers is reLU, the activation functions exiting the output layer is a sigmoid in order to approximate visit probabilities in [0, 1], and the loss function is mean squared error. The learning rate was 0.005, with a batch size of 1000, and 100 training epochs. We experimented with one, two, and three hidden layers and different numbers of hidden nodes, all with negligible effects on accuracy.

This technique is the only one of the five that needed traditional training, so we generated artificial visit subsets $V^{(s)}$ from the training data. For each value of s, we selecting 1,000,000 random subsets of size s from the training data to represent the disclosed set of visits. Each subset was paired with its corresponding full set of visits for ground truth training.

We trained a separate MLP model for each value of s. The loss function was the sum of squared errors between the binary ground truth and the [0, 1] outputs.

Testing on the same data as the other methods, the AUC values of the MLP are shown in Table I. For s = 1, MLP loses to the joint PDF and kNN. It loses slightly to kNN for s = 2 and proves superior to the other techniques for $s \ge 3$. MLP's performance improves monotonically with more disclosed locations (increasing s), as we would expect. The ROC curves for MLP are shown in Figure 5.



Fig. 4. This is the autoencoder architecture of our neural network. The drawing shows only 30 inputs and outputs, but there were actually $N_v = 3218$ of each, one for each possible visit location shown as the triangles on the maps and in Figure 1.





Fig. 5. For inferring visit locations, these are the ROC curves for the "MLP" (neural net) method, whose AUC values are given in Table I. The method performs better as the number of disclosed locations increases.

F. Summary of Location Inference

This section showed how a small location disclosure can be used to infer other places a person will likely visit. The MLP method generally worked best, achieving an AUC of 0.920 for five location disclosures. This compares to an AUC of 0.843 for the prior method, which represents inferences made without seeing any location disclosures.

Our goal is to both understand the sensitivity of location disclosures and communicate this sensitivity to normal people to help them understand the implications of a disclosure. While machine learning experts understand the AUC and F-score, they are not necessarily the best metrics to report generally, because they hide the details of the classification threshold and true/false positives. In an effort to simplify the results, we choose the operating point on the ROC curve that is nearest to the ideal operating point of (false positive rate (FPR), true positive rate (FPR)) = (0,1), i.e. the upper left corner. This is one of several different methods for choosing an operating point [22]. For the prior method, this gives (FPR,TPR) = (0.247,0.769), and it may be reasonable to report that the "accuracy" of the method is the TPR as a simplified, understandable metric, as long as the FPR is not unreasonably high. Using this approach, Table II gives the FPR and TPR for the prior method and the MLP method for different disclosure sizes |s|. We see the "accuracy" (TPR) rises with the disclosure size, as expected. With a disclosure size of five, the TPR is 8.2 percentage points above the TPR of the prior.

While reporting TPR may be effective for conveying disclosure risk, the implications of location inferences may not be clear to a regular person. In the next section we explore an implication of these types of inferences that may be more relatable.

TABLE II PERFORMANCE OF MLP METHOD FOR LOCATION INFERENCE

Dis- closure Size	False Positive Rate	True Positive Rate	TPR Improvement over Prior
0 (prior)	0.246	0.769	0.000
1	0.207	0.794	0.025
2	0.190	0.815	0.046
3	0.176	0.827	0.058
4	0.170	0.839	0.070
5	0.163	0.851	0.082

V. SENSITIVITY ANALYSIS OF POI PROPENSITY

The previous section showed how disclosing a small amount of location data can be used to infer other likely visit locations. This may be be enough to make some people take location disclosure more seriously, but there is more possible. In this section we show that small location disclosures can be used to infer the propensity of someone to visit near certain types of points of interest (POI), which may be more sensitive. For instance, it may be unnerving for someone to inadvertently reveal their high propensity to visit locations near health clinics, marijuana shops, or gambling establishments.

Toward this end, we have a list of POI and their business types covering our study area from the same supplier as our mobility data. Each POI has one of $N_b = 183$ different categories. Some of the categories are shown in Figure 6. For person *i*, the propensity to visit near a POI type *k* is $P(b_{ik} = 1)$, which is between zero and one. Recall that the binary variable v_{ij} indicates the ground truth of whether or not person *i* visited cell *j*. We also know the categories of all the POI inside each cell. This is represented by the indicator function $\mathbb{1}_{jk}$, which is one if cell *j* contains at least one POI of type *k* and zero otherwise. In our case, $k \in \{1, 2, 3, ..., N_b\}$. Thus the ground truth propensity of person *i* visiting near a POI type *k* is then

$$P^{(g)}(b_{ik} = 1) = p_{ik}^{(g)} = \frac{\sum_{j=1}^{N_b} v_{ij} \mathbb{1}_{jk}}{\sum_{j=1}^{N_b} v_{ij}}$$
(3)

In words, the visit propensity for a POI type k (e.g. restaurant) is the number of visits to a cell that contain at least one instance of the POI type normalized by the total number of

distinct cells visited. Because we do not track multiple visits to the same cell, this propensity ignores repeat visits to the same POI, including POI that might be near the person's home. Instead, it reflects visits to *different* locations that contain the same POI type. This is appropriate for an advertiser who wants to find ad recipients who are likely willing to try a new POI of a type they visit frequently, because the person already visits near the POI type in different places. Also, a visit to a geographic cell is not necessarily proof of a visit to POI in that cell. However, we make the assumption that a cell visit is indicative of a POI visit, as a way to compensate for the sporadic sampling rate of our location measurements, their spatial uncertainty, and the fact that people often walk to different POI between location measurements. We acknowledge that these propensities are also sensitive to the density of various POI types, with higher densities leading to higher propensities. This is why we look for people with significant deviations from the mean propensities, which are more likely induced by behavior rather than POI density.

Using Equation 3, we can compute the ground truth POI category visit propensity $p_{ik}^{(g)}$ for all our test users. The top three are "Restaurants and Other Eating Places", "Personal Care Services", and "Offices of Other Health Practioners".

Ultimately we would like to make relatable inferences about people to help them understand their privacy risk, such as "You tend to visit near more health clinics than most people." Thus we look for a person's POI propensities that are unusually high. For each POI type k, we use our training data to compute the mean visit propensity p_k and the standard deviation of the visit propensity σ_k . We declare an unusually high propensity if a person's normalized visit propensity exceeds a threshold $T \in [0, 1]$, i.e. if

$$\frac{p_{ik} - p_k}{\sigma_k} > T \tag{4}$$

where p_{ik} is the estimated POI propensity. This inequality is true if an individual's POI category propensity p_{ik} exceeds the POI category's mean propensity p_k by a fraction of Tof the category's propensity standard deviation σ_k . If this inequality holds for some T, then we declare that the person has an unusual affinity for places with POI type k. For our experiments, we set T = 0.5.

We next describe two different methods for inferring unusually high POI propensities.

A. Propensity from Inferred Visits

From the analysis in Section IV, we have various methods to infer the probability of a visit to a location on the map, i.e. $P(v_{ij} = 1|V_i^{(s)})$ for person *i* visiting location cell *j* given *s* previously recorded visits $V_i^{(s)}$. Adapting Equation 3, we can compute an estimated POI propensity in expectation as

$$P^{(l)}(b_{ik} = 1) = p_{ik}^{(l)} = \frac{\sum_{j=1}^{N_b} P(v_{ij} = 1 | V_i^{(s)}) \mathbb{1}_{jk}}{\sum_{j=1}^{N_b} P(v_{ij} = 1 | V_i^{(s)})}$$
(5)

Here we have just replaced v_{ij} in Equation 3 with $P(v_{ij} = 1|V_i^{(s)})$, the visit probability we estimated in Section IV. The superscript l indicates it was estimated from location visit estimates. The approach in this subsection is to first estimate the visited locations as in Section IV and then use these inferred visits to estimate the POI propensity. We use the visit probabilities from the MLP method described in Section IV-E.

Recall that the ground truth POI propensity is $p_{ik}^{(g)}$ from Equation 3, and the POI propensity estimated from the inferred visits is $p_{ik}^{(l)}$ from Equation 5. For evaluation, we could simply compare these two quantities over all persons *i* and POI types *k*. Alternatively, we could threshold both according to Equation 4 and compare the resulting booleans. Instead, we evaluate the method by first thresholding the ground truth propensities $p_{ik}^{(g)}$ according to Equation 4, with the resulting booleans serving as a ground truth binary classification giving which POI categories receive an unusally high number of nearby visits from person *i*. Then we use the inferred propensities $p_{ik}^{(l)}$ as classification probabilities, which means we can apply AUC again for evaluation.

The resulting AUC values are shown in Table III for varying sizes of location disclosures. The AUC varies from 0.740 for a disclosure of one location to 0.809 for a disclosure of five locations. Even small disclosures are revealing.

B. Propensity from Disclosed Visits

The previous subsection started with a location disclosure and used a probabilistic inference of visited cells to estimate POI category propensities. In this subsection, we skip the inference of visited cells, instead going directly from disclosed locations to high propensity POI categories. This is done with another neural net configured as an autoencoder, illustrated in Figure 6. The only difference between this network and the one in Section IV-E is that this one's output layer has a node for each POI category rather than for each location cell. Otherwise the learning parameters, activation functions and loss function are identical. The output vectors in the training phase are binary vectors indicating the state of the inequality in Equation 4, essentially giving which POI categories have an unusually high visit propensity. The input vectors are the same as in Section IV-E, and the training and testing data are split in the same way. The output of this network is a vector of propensity probabilities $P^{(d)}(b_{ik}=1)=p_{ik}^{(d)}$, where the d superscript indicates they were computed directly from the location disclosures.

The results of this more direct approach are shown in Table III. The AUC of the direct inference is consistently higher than the previous method which used an intermediate estimate of visited cells. The AUC varies from 0.828 for one location disclosure to 0.884 for five location disclosures, rising monotonically with more disclosures as we would expect. Choosing the point on the ROC curve nearest (0,1) gives the false positive and true positive rates in the last two columns of Table III. Disclosing a single location gives a true positive rate (TPR) of 0.773 (FPR = 0.295), and disclosing five locations gives a TPR of 0.790 (FPR = 0.205). Thus disclosing even a

small number of locations can still lead to accurate inferences of a person's high-propensity POI categories.

Unlike the previous Section IV on location inference, this section does not compare the inferential accuracy to a prior. This is because the goal in this section is to find POI categories for which a user has an unusual propensity to be near. A prior would simply say that each user's propensities are exactly ordinary and none are unusual.



Fig. 6. This is the autoencoder architecture of our neural network for inferring unusually high POI category visits directly from disclosed locations. The drawing shows only 30 inputs and outputs, but there were actually $N_v = 3218$ binary inputs, one for each possible visit location shown as the triangles on the maps and in Figure 1. There were $N_p = 183$ binary outputs, one for each POI type.



Fig. 7. For inferring high propensity POI categories, these are the ROC curves for the "MLP" (neural net) method, whose AUC values are given in Table III. The method performs better as the number of disclosed locations increases.

VI. LOCATION BASED ADVERTISING

One consequence of disclosing location is advertising. A 2014 study of personal data brokers found that advertising is the primary use of personal data [2]. Mobile ads can be effective: one study saw ad recipients respond with a purchase between 2.1% and 4.3% of the time [23]. This section presents a simplified model of how an advertiser could use inferences about POI category preferences from the previous section to

TABLE III AUC values for POI category visits

Dis- closure Size	Inferred Location \rightarrow POI AUC	Disclosure →POI AUC	Disclosure →POI FPR	Disclosure →POI TPR
1	0.740	0.828	0.295	0.773
2	0.752	0.849	0.259	0.771
3	0.630	0.863	0.242	0.780
4	0.653	0.876	0.221	0.784
5	0.809	0.884	0.205	0.790

TABLE IV Payoff matrix for ad delivery

		POI Category Propensity	
		no propensity	high propensity
Ad	do not deliver	$b_{11} = 0$	$b_{12} = -1.0 + \alpha$
	deliver	$b_{21} = -\alpha$	$b_{22} = 1.0 - \alpha$

decide whether or not to deliver an ad. The decision flows naturally from the POI propensity ROC curves (Figure 7) and a payoff matrix that gives the economic return of delivering or not delivering an ad, highlighting how a location disclosure can help an advertiser.

The assumed payoff matrix is shown in Table IV. The advertiser uses this to express the amount of gain or loss depending on the choice of delivering or not delivering an advertisement to someone who might show a high propensity for visiting a certain category of POI. For instance, a restaurant chain might use to this for delivering ads to people who show an unusually high propensity for visiting locations with restaurants, as described in Section V. This is similar to the payoff matrix from Aly et al. [24].

The first row of the payoff matrix gives the consequences for not delivering an ad. In the first column of the first row, the payoff for not delivering an ad to someone with no POI propensity is b_{11} in general. This is the correct action to take in this case, and the payoff is $b_{11} = 0$, because there is no ad cost and no benefit. The lower left element gives the cost of delivering an ad to someone with no POI propensity. This is simply the cost cost of the ad, α , in some monetary units. For instance, this would cover the case of delivering a restaurant ad to someone who does not show any unusual propensity to visit locations with restaurants. The lower right element of the payoff matrix is b_{22} , and it gives the benefit of delivering an ad to someone with a high POI propensity. This is the right action from the advertiser's standpoint. Because we lack the specifics of the POI type, responses rates, and profit margins, we say the expected benefit of a proper ad is 1.0, which is primarily the expected profit from a welltargeted ad. We subtract the cost of the ad, making the payoff $b_{22} = 1 - \alpha$. Because we assume an expected profit of 1.0, all the other payoff values are implicitly expressed as proportions of this payoff for our model. The upper right element, b_{12} , gives the cost of not delivering an ad to a high propensity individual. This is an opportunity cost, because the person should have received the ad, but did not. We approximate the

cost of this lost opportunity as the negative of the benefit of properly delivering an ad. Thus $b_{12} = -b_{22} = -1.0 + \alpha$. The α of this term can be considered the savings associated with not buying an ad. The goal of this analysis will be to see how much an advertiser can afford to pay for an ad and still maintain a positive return, given the quality of inferences possible from Section V on POI propensity.

The expected return on advertising is a function of the quality of the POI propensity inferences and the payoff matrix. Each cell of the payoff matrix corresponds to one measure of inference quality. For instance, the lower right element, properly delivering an ad to a high propensity person, corresponds to the true positive rate of the POI propensity classifier described in Section V. Using the true positive rate (TPR), false positive rate (FPR), true negative rate (TNR), and false negative rate (FNR), the expected payoff of an advertising campaign, per individual, is

$$\mathbb{E}[P] = b_{11} \cdot \text{TNR} + b_{12} \cdot \text{FNR} + b_{21} \cdot \text{FPR} + b_{22} \cdot \text{TPR} \quad (6)$$

We know that TNR = 1 - FPR and FNR = 1 - TPR. Using these substitutions along with the modeled values of the payoff matrix from Table IV, we have

$$\mathbb{E}[P] = (1 - 2 \cdot \text{TPR} - \text{FPR})\alpha + 2 \cdot \text{TPR} - 1$$
(7)

The advertiser wants a positive return, i.e. $\mathbb{E}[P] > 0$, which implies

$$\alpha < \frac{1 - 2 \cdot \text{TPR}}{1 - 2 \cdot \text{TPR} - \text{FPR}} \tag{8}$$

This gives an upper bound that an advertiser would be willing to pay for an ad while still maintaining an overall positive return on the advertising campaign. If this value is fairly high, then the advertising campaign would be more attractive, resulting in a higher chance of the individual's location disclosure triggering an ad.

Figure 7 gives the FPR and TPR values for inferring POI propensity from the neural net model. These values represent the classifier quality that the advertiser is constrained to work with, although the advertiser may choose any (FPR,TPR) pair along the curves for flexibility. Inserting these values into Equation 8, we see the maximum tolerable values for α in Figure 8.

Figure 8 shows several interesting characteristics. Recall that the horizontal axis gives the operating point of the POI propensity classifier in terms of TPR. From Table III, the TPR of the ROC nearest (0,1) is about 0.780 for all five disclosure sizes. From the plot in Figure 8, this translated to about $0.65 < \alpha < 0.75$, which represents the maximum tolerable cost of an ad. The benefit of properly delivering an ad was set at 1.0, not counting the cost of the ad. The plot is indicating that an ad can tolerably cost up to 65% to 75% of the benefit of a useful ad for our particular choice of parameters. The plot also shows that at a TPR of less than 0.5, the tolerable ad cost is negative, meaning the ad campaign would likely be

Maximum Tolerable Cost of Ad



Fig. 8. For each location disclosure size s, this gives the maximum relative amount α that an advertiser could tolerate to pay for an ad. The best true positive rate (TPR) for the POI propensity inferences is about 0.780 (Table III), meaning that the tolerable cost is around $0.65 < \alpha < 0.75$ for all five disclosure sizes.

unprofitable. Finally, as expected, the tolerable cost of an ad rises with the size of the location disclosure *s*, because the POI propensity classifier becomes more accurate.

This model shows how to quantify the costs and benefits of delivering ads based on location disclosures, making it easier to anticipate the consequences of an individual revealing even a few location visits. With specific profit estimates and ad success rates for different POI types, the model can show whether or not a location disclosure might trigger an ad.

VII. SUMMARY AND CONCLUSION

This work addresses a new problem: analyzing the sensitivity of disclosing a few personal locations from the individual's point of view. This is an important component of understanding the privacy implications of sharing location data with an enterprise whose inferential "black boxes" are proprietary. We showed how a small number of location disclosures can be used to infer other likely visit locations, and we quantified the accuracy boost over a prior distribution. Personal location data reveals a person's propensity to visit places with certain types of businesses, and our analysis showed how to discover those types for which the user seems to have an unusually high propensity. One of the major uses of personal location data is advertising, so we concluded with an economic analysis of ad delivery based on inferred propensities for certain business types. For all our analyses, we tested and quantified the inferential accuracy using personal location data from a typical data broker.

This work is a first step toward understanding and communicating the consequences of disclosing personal location data for regular people, which is helpful for making informed decisions about location privacy. Future work should explore more sophisticated inferences as well as the effects of privacy techniques like obscuring location data with random noise.

References

- Valiator, "Privacy vs. personalisation," https://www.valitor.com/ wp-content/uploads/2019/11/apex-privacy-report-nov-2019-low.pdf, Tech. Rep., 2019, accessed: 30-May-2021.
- [2] E. Ramirez, J. Brill, M. K. Ohlhausen, J. D. Write, and T. McSweeny, "Data brokers: A call for transparency and accountability," U.S. Federal Trade Commission, Washington, DC, Tech. Rep., May 2014.
- [3] J. Valentino-DeVries, N. Singer, M. H. Keller, and A. Krolik, "Your apps know where you were last night, and they're not keeping it secret," *The New York Times*, [Online; accessed 30-May-2021]. [Online]. Available: https://www.nytimes.com/interactive/2018/12/10/ business/location-data-privacy-apps.html
- [4] S. Wachter, B. Mittelstadt, and L. Floridi, "Why a right to explanation of automated decision-making does not exist in the general data protection regulation," *International Data Privacy Law*, vol. 7, no. 2, pp. 76–99, 2017.
- [5] N. Gerber, P. Gerber, and M. Volkamer, "Explaining the privacy paradox: A systematic review of literature investigating privacy attitude and behavior," *Computers & security*, vol. 77, pp. 226–261, 2018.
- [6] E. Hargittai and A. Marwick, ""what can i really do?" explaining the privacy paradox with online apathy," *International journal of communication*, vol. 10, p. 21, 2016.
- [7] S. Kokolakis, "Privacy attitudes and privacy behaviour: A review of current research on the privacy paradox phenomenon," *Computers & security*, vol. 64, pp. 122–134, 2017.
- [8] A. M. Zafeiropoulou, D. E. Millard, C. Webber, and K. O'Hara, "Unpicking the privacy paradox: can structuration theory help to explain location-based privacy decisions?" in *Proceedings of the 5th Annual* ACM Web Science Conference, 2013, pp. 463–472.
- [9] H. Almuhimedi, "Helping smartphone users manage their privacy through nudges," Ph.D. dissertation, Carnegie Mellon University, Pittsburgh, PA, December 2017.
- [10] N. Marmasse and C. Schmandt, "A user-centered location model," *Personal and ubiquitous computing*, vol. 6, no. 5, pp. 318–321, 2002.
- [11] D. Ashbrook and T. Starner, "Using gps to learn significant locations and predict movement across multiple users," *Personal and Ubiquitous computing*, vol. 7, no. 5, pp. 275–286, 2003.
- [12] C. Song, Z. Qu, N. Blumm, and A.-L. Barabási, "Limits of predictability in human mobility," *Science*, vol. 327, no. 5968, pp. 1018–1021, 2010.
- [13] N. C. Chen, W. Xie, R. E. Welsch, K. Larson, and J. Xie, "Comprehensive predictions of tourists' next visit location based on call detail records using machine learning and deep learning methods," in 2017 *IEEE International Congress on Big Data (BigData Congress)*. IEEE, 2017, pp. 1–6.
- [14] W. Zheng, X. Huang, and Y. Li, "Understanding the tourist mobility using gps: Where is the next place?" *Tourism Management*, vol. 59, pp. 267–280, 2017.
- [15] C. I. Muntean, F. M. Nardini, F. Silvestri, and R. Baraglia, "On learning prediction models for tourists paths," *ACM Transactions on Intelligent Systems and Technology (TIST)*, vol. 7, no. 1, pp. 1–34, 2015.
- [16] Q. Song, H. Ge, J. Caverlee, and X. Hu, "Tensor completion algorithms in big data analytics," ACM Transactions on Knowledge Discovery from Data (TKDD), vol. 13, no. 1, pp. 1–48, 2019.
- [17] J. Bobadilla, F. Ortega, A. Hernando, and A. Gutiérrez, "Recommender systems survey," *Knowledge-based systems*, vol. 46, pp. 109–132, 2013.
- [18] Wikipedia contributors, "Netflix prize Wikipedia, the free encyclopedia," https://en.wikipedia.org/w/index.php?title=Netflix_Prize&oldid= 1020767219, 2021, [Online; accessed 16-May-2021].
- [19] J. Froehlich, M. Y. Chen, I. E. Smith, and F. Potter, "Voting with your feet: An investigative study of the relationship between place visit behavior and preference," in *International Conference on Ubiquitous Computing.* Springer, 2006, pp. 333–350.
- [20] M. Bertalmio, G. Sapiro, V. Caselles, and C. Ballester, "Image inpainting," in *Proceedings of the 27th annual conference on Computer* graphics and interactive techniques, 2000, pp. 417–424.
- [21] A. S. Szalay, J. Gray, G. Fekete, P. Z. Kunszt, P. Kukol, and A. Thakar, "Indexing the sphere with the hierarchical triangular mesh," *arXiv* preprint cs/0701164, 2007.
- [22] F. Habibzadeh, P. Habibzadeh, and M. Yadollahie, "On determining the most appropriate test cut-off value: the case of tests with continuous results," *Biochemia medica*, vol. 26, no. 3, pp. 297–307, 2016.

- [23] M. Andrews, X. Luo, Z. Fang, and A. Ghose, "Mobile ad effectiveness: Hyper-contextual targeting with crowdedness," *Marketing Science*, vol. 35, no. 2, pp. 218–233, 2016.
- [24] H. Aly, J. Krumm, G. Ranade, and E. Horvitz, "On the value of spatiotemporal information: Principles and scenarios," in *Proceedings* of the 26th ACM SIGSPATIAL International Conference on Advances in Geographic Information Systems, 2018, pp. 179–188.