Linking Anonymous Location Traces Through Driving Characteristics

Bin Zan
WINLAB, Rutgers University
671 Route 1 South
North Brunswick, NJ
08902-3390
zanb@winlab.rutgers.edu

Zhanbo Sun
Rensselaer Polytechnic Institute
110 8th St
Troy, NY 12180-3590
sunz2@rpi.edu

Marco Gruteser
WINLAB, Rutgers University
671 Route 1 South
North Brunswick, NJ
08902-3390
gruteser@winlab.rutgers.edu

Xuegang Ban
Rensselaer Polytechnic Institute
110 8th St
Troy, NY 12180-3590
banx@rpi.edu

ABSTRACT

Efforts to anonymize collections of location traces have often sought to reduce re-identification risks by dividing longer traces into multiple shorter, unlinkable segments. To ensure unlinkability, these algorithms delete parts from each location trace in areas where multiple traces converge, so that it is difficult to predict the movements of any one subject within this area and identify which follow-on trace segments belongs to the same subject. In this paper, we ask whether it is sufficient to base the definition of unlinkability on movement prediction models or whether the revealed trace segments themselves contain a fingerprint of the data subject that can be used to link segments and ultimately recover private information. To this end, we study a large set of vehicle locations traces collected through the Next Generation Simulation program. We first show that using vehicle moving characteristics related features, it is possible to identify outliers such as trucks or motorcycles from general passenger automobiles. We then show that even in a dataset containing similar passenger automobiles only, it is possible to use outlier driving behaviors to link a fraction of the vehicle trips. These results show that the definition of unlinkability may have to be extended for very precise location traces.

Categories and Subject Descriptors

K.4.1 [**COMPUTERS AND SOCIETY**]: Public Policy Issues—*Privacy*

Keywords

privacy, mix-zone, anonymity, outlier, ROC curve

CODASPY'13, February 18–20, 2013, San Antonio, Texas, USA.

1. INTRODUCTION

The broad adoption of location-based services has led to an increasing number of services and applications that monitor and record time-series location traces of peoples movements [16, 19, 18, 22]. While some applications require knowledge of the user, a number of applications can enhance privacy by working with anonymous location records. Traffic engineering-related applications that monitor traffic states and other performance measures are one example of such applications.

As in other contexts, achieving strong anonymity in a dataset of location records requires more effort than just deleting identifiers such as user ids, login names, cell phone IDs, or IP addresses [7, 17, 4]. For this reason, a frequently proposed technique to enhance the anonymity of location records is to delete portions of the location traces themselves. This includes trip starting/ending points, which might point to a sensitive location where a user does not want to be seen or a particularly identifying location such as a home, which would make it easier to re-identify a location trace [11, 3].

Since it is difficult, however, to scrub all such locations from traces, a frequently proposed concept is to also delete portions of the data within trips, so that longer trips are divided into shorter unlinkable segments. This will limit the privacy leakage to the information within one short segment, if a re-identification at one location were to occur. Variations of this concept are known under the names of mix-zones [5, 9] or path cloaking [13, 23], and have been the subject of many follow-up studies (e.g., [6, 20, 15, 10, 21]). These approaches do have in common that they define and evaluate unlinkability primarily through the use of movement prediction models. Informally, two trip segments are unlinkable if it is difficult to predict the missing portion of the path with sufficient precision to determine that these two segments are indeed recorded from the same user.

In this paper, we test whether this assumption remains sufficient as the time-series location traces become more precise. For example, typical GPS receivers in vehicles can provide location samples at an updating rate of 1 Hz and with an error of less than 3 m in most cases. Given such precise movement traces, are there other unique patterns embedded in the traces, which would allow linking of two location trace segments without using the prediction models?

One form of such unique patterns are outliers in vehicle moving

characteristics. We studied a large set of vehicle traces and tried to identify such outliers. In particular, we considered an outlier as a vehicle which exhibits higher speed, higher acceleration, or a larger number of lane changes compared to the other nearby vehicles, which are grouped inside the same anonymity group. Here, an anonymity group can be considered as a number of close-by vehicles whose trace data are available to the adversary without direct associated information to their real IDs.

The reasons a vehicle may have very distinguishable moving characteristics comparing to others can be classified into two categories: intrinsic and extrinsic. Intrinsic reasons are due to the physical nature of a vehicle, for example a large and heavy delivery truck usually has smaller mean acceleration than a small passenger automobile. Extrinsic reasons are the result of the driver's driving propensity/behavior or other external conditions. For example, many of us tend to drive faster when we are pressed by time. The intrinsic and extrinsic causes are not absolutely disconnected from each other. Many times they are related and both contribute to distinguishable moving characteristics.

The rest of the paper is organized as follows. In section 2, we briefly describe the adversary model. In section 3, we discuss our motivation with a set of data analysis results. Section 4, we study several possible learning models an adversary can use to attack existing mix-zone models. In section 5, we evaluate the proposed learning models and the impacts on the mix-zone model. Section 6 concludes our work.

2. ADVERSARY MODEL

We consider a scenario where an adversary, perhaps an insider, has gained access to a database with the location trace information from drivers. A location trace is defined by a sequence of timestamped location records with approximately meter-level precision at a time resolution of about one second. We assume that mix zone privacy techniques have been applied: all identifiers such as vehicle license plates associated with the traces have been removed. To distinguish different traces in this dataset, they have been replaced with pseudo identifiers. Furthermore, a long trace is split into several short segments with different pseudo-IDs in order to reduce the possible information leakage due to long term tracking [3]. The mix-zone model [5] is assumed to be used in which all the location traces are discarded inside mix-zones, leaving only segments outside that area. Based on all the available trace segments (which may belong to different users but are all mixed together), an adversary's task is to identify which segments were generated by the same user. In other words, the adversary tries to link segments after the mix zone to a segment before the mix zone (which the adversary may have already correctly associated with the target).

3. FEASIBILITY OF LOCATION TRACE OUTLIER DETECTION

So far the analysis on privacy preserving models has been purely based on predicting and matching the reemergence of vehicles out of the mix-zone. That is, the adversary can link two trace segments if he/she can correctly predict where or when a vehicle will appear. The privacy results of such an analysis are based on the hidden assumption that no other way of linking vehicle traces exists. In this section, we will explore the feasibility of linking trace segments based on characteristics of their movement.

3.1 The Rise of an Outlier

Consider the example shown in Fig. 1 where one vehicle tends to drive significantly faster than other nearby vehicles.

Figure 1: A fast car appears as an outlier.

Let us assume that all the vehicle traces outside the mix-zone are available to the adversary since vehicle a has very different movement characteristics (higher speed) both before the mix-zone and after, the adversary could assume that these two trace segments were generated by the same vehicle and link them into a single trace. In practice, the success of such heuristics will depend on actual speed distributions as well as their tendency to maintain speed long enough to traverse a mix-zone. We refer to vehicle a an outlier because its movement pattern is very different from the others. The above example actually indicates that outliers could destroy the mix-zone mechanism for the easy linkability between the traces.

In [12], Grubbs defined an outlier as: "one that appears to deviate markedly from other members of the sample in which it occurs." When considering all trace segments of a group of vehicles as samples, an outlier segment must appear to deviate markedly from the others. As discussed above, special movement characteristics could lead to an outlier. The cause of such special characteristics can be considered from two sides: intrinsic and extrinsic. The difference between a typical delivery truck and a typical passenger automobile is intrinsic (similarly between a motorcycle and an automobile). On the other hand, a particular driver's speed as mentioned in the above example is an extrinsic cause of being an outlier. Our work will focus on intrinsic reasons first and then move to the extrinsic causes.

According to the US Bureau of Transportation Statistics [1], there were 8,212,267 motorcycles, 10,770,054 trucks and 230,444,440 passenger automobiles in the US in 2010. Therefore, we will see only one truck for every 25 vehicles observed on average (Note that the ratio of trucks and motorcycles is only 4.4% and 3.3% of the total number of vehicles, respectively). From everyday observation, trucks and motorcycles have different movement patterns from passenger automobiles. An adversary can exploit this knowledge to help him/her to identify a truck (or a motorcycle) target from a group of automobiles and to link the trace segments of the truck (or motorcycle) before and after mix-zones. Considering the ratio of trucks and motorcycles to the total number of vehicles on the road, such a scenario will stand out and be easily observed.

3.2 The Observation of Real World Driving Characteristics

In this subsection, we use a set of results to illustrate the feasibility of outlier detection based on vehicle movement characteristics. Since extrinsic cases are based more on individual behavior, we will focus first on discussing cases that are based on intrinsic reasons. The following results are based on NGSIM [2] data which consists of detailed vehicle trajectories, wide-area detectors, and supporting data from researching driver behavior. The vehicle trajectory data was collected using digital video cameras to record the precise location of each vehicle on a 0.5- to 1.0-kilometer section of roadway every one-tenth of a second. The portion of data we use covers the southbound direction of highway US 101 (Hollywood Freeway) in Los Angeles, California. The video cameras were mounted on a 36-story building, 10 Universal City Plaza, which is located adjacent to the U.S. Highway 101 and Lankershim Boulevard interchange in the Universal City neighborhood.

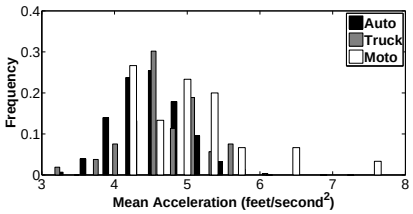

Figure 2: Histogram of mean acceleration for trucks, autos and motos.

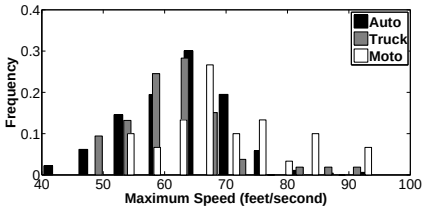

Figure 3: Histogram of maximum speed for trucks, autos and motos.

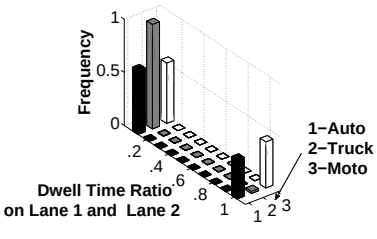

Figure 4: Histogram of dwell time ratios on left most lanes for trucks, autos and motos.

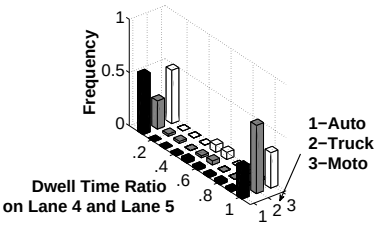

Figure 5: Histogram of dwell time ratios on right most lanes for trucks, autos and motos.

The vehicle trajectory data was transcribed from the video data using a customized software application developed for NGSIM. The total length of the road segment is about 2100 feet and there are five main lanes throughout the section. Lane numbers start from the left-most lane. The total time spans from 7:50am to 8:05am on June 15, 2005. There are 2086 passenger automobiles (autos), 53 trucks and 30 motorcycles (motos) traces in the dataset. Note, the distribution of vehicle types in this dataset is similar to the US Bureau of Transportation Statistics [1].

First, we show a histogram plot of mean acceleration for motos, trucks and autos in Fig. 2. As can be seen, motorcycles tend to have a larger mean acceleration compared to automobiles and trucks (e.g., most of the motorcycles show a mean acceleration higher than 4.2 ft/s^2). The acceleration of trucks tends to be smaller compared to other vehicle types. As shown in the figure, the minimum mean acceleration belongs to trucks. Under certain circumstances (e.g., the target is the only typical motorcycle among a group of automobiles), an adversary could easily identify the target solely based on the difference in the mean acceleration between these two types of vehicles. This ability to distinguish between vehicles increases the linkability of vehicle traces before and after mix-zones.

Similar observations can also be seen in other vehicle movement characteristics. Fig. 3 shows the histogram plots of the maximum speed for trucks, automobiles and motorcycles. As with the mean acceleration, most motorcycles tend to have larger values of maximum speed while automobiles have both larger and smaller values than trucks. The former is quite consistent with our intuition while the latter is not. To explain why automobiles have both larger and smaller maximum speeds, we note that the drivers of automobiles cover a large range of the population, so some of them tend to drive fast while many others tend to drive slowly (e.g., seniors). On the other hand, the people who drive trucks are mostly professional drivers that are either young or in their middle ages. Thus, it is also reasonable to see that the use of automobiles covers a larger range in terms of the maximum speed. However, from the figure, we can conclude again that under certain circumstances, it is easy for an adversary to identify a target vehicle if he/she can exploit the maximum speed information (e.g., when the target is a typical motorcycle and the other vehicles are all common automobiles or trucks).

The above two examples do not have a clear way to distinguish a truck from other types of vehicles based on one particular movement characteristic. However, some trends can at least be observed. For example, trucks generally have a relatively small mean acceleration and maximum speed. Our proposed adversary model does not require a strong difference in one particular movement characteristic, instead, small differences in multiple characteristics can allow for a way to distinguish target types of vehicles from the groups of vehicles.

In the next subsection, we will illustrate more possible movement characteristics that can be used for intrinsic cause outlier detection.

3.3 Exploiting Other Features

There are a few movement characteristics that may be useful for intrinsic cause outlier detection. In this subsection, we illustrate two other characteristics that are useful for identifying trucks from other vehicles.

3.3.1 Lane Changing

As more and more detailed and accurate location information becomes available, some movement characteristics (patterns) that were not available before, are easier to obtained. For example, when the GPS or other location devices become accurate enough, recognizing the lane a vehicle is moving on becomes feasible. For example, in the NGSIM [2] dataset all vehicle traces are recorded with lane information. As another example, High Accuracy Nationwide Differential GPS (HA-NDGPS) system which is currently under development can provide 10-15 centimeter accuracy.

In Fig. 4, we show the histogram of dwll time ratios on left most lanes for trucks, automobiles and motorcycles. The question of which lanes to stay may appear as a driver's preference or an extrinsic reason at the first glance. However, if the special movement

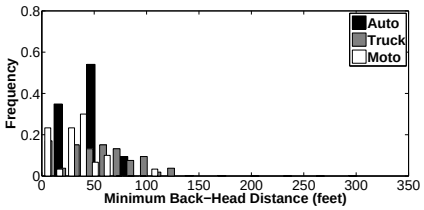

Figure 6: Histogram of minimum back-headway distance for trucks, autos and motos.

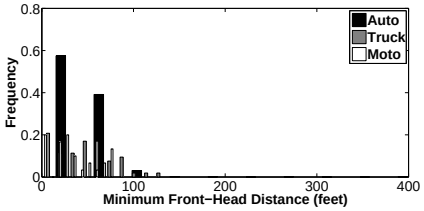

Figure 7: Histogram of minimum front-headway distance for trucks, autos and motos.

pattern appears in most trucks, we may presume that this is the result from intrinsic or mixed of the intrinsic and extrinsic reasons. In fact, it is reasonable to assume truck drivers tend to drive on the outer (right most) lanes instead of inner (left most) lanes because trucks can not move as flexible as small automobiles generally. Inner lanes are usually reserved for fast moving vehicles, slow in lane changing and speed changing discourage truck drivers to stay in the inner lanes. The result in Fig. 5 complement what is shown in Fig. 4. Trucks do not stay on the inner lanes most times, instead they tend to appear in the outer lanes. Above plots indicate the possibility for an adversary to detect a truck outlier out of automobiles and other type of vehicles.

3.3.2 Headway Distance

Some features or characteristics may not be always available or only available for some vehicles. However, once available, they may increase an adversary's ability to detect certain types of vehicles or a special kind of target. One example is the headway information. Headway is defined as the head-to-head distance between a vehicle and its immediate frontward vehicle (front-headway) or the vehicle immediate afterward (back-headway). The maximum headway distance does not help the adversary much since the values can be very large when there is no vehicle nearby. Intuitively, headway distance (especially front-headway) can be useful for distinguishing trucks and small automobiles, because trucks tend to leave more space in front of them. However, the effectiveness of using headway distance is strongly related to the penetration rate. If an adversary has obtained all vehicles' location traces, then it can compute the accurate headway distance for each one, otherwise, the value may not be correct. Consider a scenario, where vehicle a is immediately followed by vehicle 'b'. The location trace of a is not known by the adversary, then the front-headway distance of vehicle 'b' can't be computed. Therefore, the usage of headway distance is limited.

As shown in Fig. 6 and Fig. 7, trucks have the largest front-headway distance and back-headway distance compared to automobiles, while motorcycles have the smallest values. Although, this is not true for all the cases, the characteristic over large amount of samples will be useful when adding such knowledge into a machine learning classification model. As we will show in the next section, one characteristic alone may not be able to differentiate all trucks or motorcycles from automobiles, however it can be very helpful when combined with other movement characteristics.

4. TRACE OUTLIER DETECTION ALGORITHMS

The previous section shows the feasibility of location trace outlier detection. In this section, we provide detailed methods for such detection.

4.1 Intrinsic Outlier Detection

Recall that intrinsic outlier movements are caused by less common types of vehicle. In particular, we consider the case where an adversary knows a priori that the vehicle of interest is a less common vehicle (e.g., a truck or motorcycle) when trying to track the vehicle. This task would become significantly easier, if the adversary can first filter out the traces from all other vehicle types because this would result in a smaller anonymity set and less probability of linking error.

To realize this type of vehicle filtering, we use machine learning classification approaches. An adversary uses historical data to train the classification model, and then uses the model to classify and filter vehicles crossing a mix-zone. Compared to the extrinsic techniques we will discuss later, a key advantage of this method is that the training dataset is not limited to traces from the target vehicle but can also use traces from other vehicles of the same type. This results in a larger sample pool and can build a better model for outlier detection. However, the adversary must know the intrinsic nature of the target vehicle (e.g., vehicle type) and the characteristic must belong to a less common vehicle. If the confidence level is low, an adversary still has to resort to more general outlier detection strategies, to be discussed later.

In particular, we studied three classic machine learning models, and compared their performance. To be more specific, we studied Linear Discriminant Analysis (LDA), Quadratic Discriminant Analysis (QDA) and Naive Bayes Models. LDA generates linear boundaries, QDA and NB can both generate quadratic boundaries. For more details about these learning models, the interested reader could consult [14].

The input of each machine learning model are the road segment traces obtained from NGSIM data [2]. Available information includes vehicle ID, global time, local X, local Y, vehicle length, vehicle width, vehicle velocity, vehicle acceleration, lane identification, preceding vehicle, following vehicle, space headway and vehicle class, etc. Local X and Y is the distance from the front center of the vehicle to the left most edge of the road segment and to the entry edge of the segment respectively. The space headway information only covers the front-headway distance, however, with the preceding vehicle and following vehicle information, back-headway distance can be computed and used in the adversary model. For more detailed information regarding the NGSIM data, please refer to [2].

Based on above dataset, we further extract the following sample based movement characteristics:

1. speed (5 features): maximum speed, minimum speed, average speed, median speed, standard deviation of speed.
2. acceleration (8 features): maximum acceleration, maximum deceleration, average acceleration, average deceleration, median acceleration, median deceleration, standard deviation of acceleration, standard deviation of deceleration.
3. proportion (4 features): acceleration greater than 5 ft/s^2, de-

celeration greater than 5 ft/s^2, speed greater than 60 ft/s, speed less than 20 ft/s.

4. lane position (6 features): frequency on visiting lane 1, 2, 3, 4, 5, total number of lane switching.
5. headway (4 features): minimum distance to the vehicle in front, average distance to the vehicle in front, minimum distance to the vehicle after, average distance to the vehicle after.

Before illustrating our proposed adversary models in the next subsection, we go through several initial observations from studying.

- Some movement characteristics may help an adversary more than others with building the outlier detection models. For example, in the 1927 data samples collected from US101 highway between 7:50am and 8:05am (including both automobiles and trucks), the residual error of a QDA model on maximum speed is only 0.34. While with minimum speed, the error rate is 0.56 which is much larger.
- A movement characteristic itself may not be good enough to differentiate vehicles. However, a combination of such characteristics can be useful. For example, using LDA to distinguish three types of vehicles (trucks, automobiles and motorcycles) from the same dataset with average acceleration, the (residual) error rate is 0.41. If based on the combination of multiple dimensions: maximum speed, average acceleration, proportion of deceleration, the ratio of frequency of visiting lane 4 and 5, and the number of lane changes, the error rate reduce to 0.11.
- Increasing the total number of movement characteristics to be used may not always lead to better performance.
- To find a set of boundaries for classifying trucks, automobiles and motorcycles through one learning model is not easy. Instead, we choose to train a set of models each of which can distinguish a pair of vehicle types.

4.2 Feature Selection or Dimension Reduction

For building proper learning models, it is necessary to filter some characteristics (or features) which cause more noise than they contribute to the outlier detection. Therefore, we studied two methods to improve the performance.

4.2.1 Manual Feature Selection

The first strategy is to select a number of features which performs well individually. Multiple selected features are combined together to form learning models in high dimensions. For example, assume we extract a total of n features from the vehicle trace dataset, then for each feature, a learning model is built. All features are then ordered based on the performance in training dataset. Next, the top m features are selected to form the best combination. This method has several advantages. First, the resulted model is easier to be understood. Each feature corresponds to a movement characteristic which has clear definition. Second, the performance can be estimated. Since all the features have real physical meaning, it is relative easy to estimate the performance based on user's daily observation. One disadvantage is that such a method is slow. For every pair of features, the adversary needs to generate a learning model, and sort the results based on performance. Another disadvantage is that the result may not be optimal. Some features may have correlation, thus putting them together may not contribute more information for a learning model.

As part of our evaluation results, the basic features we manually selected are maximum speed, average acceleration, proportion of deceleration greater than 5 ft/s^2, frequency on visiting lane 4 and 5.

Figure 8: A general outlier detection scenario

4.2.2 Principle Component Analysis Based Dimension Reduction

The second strategy is to use a dimension reduction method. In our study, we use Principle Component Analysis (PCA) to project all the features onto m best dimensions and then train the learning model based on these m dimensions. The advantage of this method is that it is fully automatic and it can achieve better performance than the previous method. That is because if two features are strongly correlated, it will only appear as one dimension after projection in PCA. One disadvantage of such a method is the selected dimensions in the end may not have clear physical meaning, thus the results cannot be always interpreted easily.

4.3 General Outlier Detection: Extrinsic

In the extrinsic case, it is not clear that we can build a learning model based on historical data from different drivers and different trips. For example, the special movement characteristics of a driver during a trip (e.g., due to time pressure) may not appear in his/her previous trip two days ago. And there is no training dataset available from general location trace data to indicate if a driver is under time pressure or not. Comparing to intrinsic, extrinsic is more unstable and more case dependent and/or time dependent. Therefore, the machine learning model for outlier detection must be built on the fly and in real time. Assume that before a mix-zone, the trajectory of a target vehicle a is known as shown in Fig. 8, the challenge for an adversary is now to identify the same target vehicle after the mix-zone. Since there is no historical data that can be used, the size of the training dataset relies on the length of the road segment before the mix-zone. While we still can use machine learning models, we need assume the target vehicle itself is a class. There are two strategies available for the adversary:

4.3.1 One-to-One

In this method, the trace from the target vehicle forms a class, and all the other traces from the remaining vehicles of the anonymity set form the second class. The learning model detects the target based on binary classification result for the dataset after the mix-zone. This method is easy to implement, however due to the small number of data samples available in the first class compared to the second class, the classification model may not be able to generate a good boundary for a binary classification model.

4.3.2 One-to-Many

In this method, not only the trace from the target vehicle forms a class, but also all other vehicles form classes independently. The adversary builds multiple binary classification learning models for the target with all the other vehicles. To detect the outlier, the adversary runs all learning models for each trace segment collected after mix-zone to determine if a trace looks more like the target or one of the others.

5. EVALUATION

In this section, we study the proposed outlier detection techniques with the NGSIM [2] dataset. The True Positive Rate (TPR) and the False Positive Rate (FPR) are used as measures to characterize the ability of an adversary to track a vehicle through a mix-zone. The same dataset which has been studied in the section 3.2 is used. Here we give a brief introduction of the dataset again. All the location traces are collected from the US101 highway. The dataset covers the traffic data between 7:50am and 8:05am. It includes a total of 1875 automobiles, 52 trucks and 25 motorcycles, all of which have trajectories longer than 2000 feet. To date, we have focused more on intrinsic factors, since there is a relative large training dataset available. A small number of vehicle traces are removed. These are those which 1) were less than 2000 feet long; 2) never had any vehicle driving in front; 3) never had any vehicle following behind. The first condition is to reduce the possible impact from factors other than the vehicle movement characteristics, the second and third conditions are to make sure we can use the same dataset when comparing the performance of different outlier detection algorithms (with and without headway information).

The results presented here focus on showing the Receiver Operating Characteristic (ROC) curves of different learning models. An ROC graph is a technique for visualizing, organizing and selecting classifiers based on their performance. It is not only a generally useful performance graphing method, but also very useful for domains with skewed class distribution and unequal classification error cost [8]. The data source input into the outlier detection algorithm mostly follows skewed class distributions. Thus, it is beneficial to use the ROC curves for comparing the performance of different learning models. On the other hand, through the ROC curves, the tradeoff between quantity of attacks and quality of attacks (in terms of confidence level) can be easily observed.

5.1 Outlier Detection: Intrinsic

First, we show the results based on manual feature selection. The performance of the three different machine learning models LDA, QDA and Naive Bayes for truck detection are shown in Fig. 9. The main point is that by properly selecting a learning model an adversary can identify a truck or a motorcycle from common automobiles. This means that additional information leaks from the traces that could be used to compromise mix-zone protection. The features we manually selected are maximum speed, average acceleration, proportion of deceleration greater than 5 ft/s^2 , frequency of visiting lane 4 and 5. As shown in the ROC graphs, in general, to improve the confidence of tracking a target, an adversary has to sacrifice the quantity of the overall attacks. As the confidence level increases, the number of vehicles that can be tracked decreases. Nonetheless, even an adversary who can claim a very high confidence in tracking a target just occasionally may be unacceptable.

Fig. 9(a), Fig. 9(b) and Fig. 9(c) show results (for the LDA, QDA and Naive Bayes models) from the training dataset and results of 10-fold cross validation on the QDA as well as Naive Bayes model, respectively. As can be noticed from Fig. 9(a), all the three learning models generate better ROC curves than randomly guessing (the diagonal line in the figure) in training. This indicates that with machine learning classification models an adversary is able to identify trucks from automobiles under certain circumstances. To be more specific, the ROC curve of the QDA model crossing the point at FPR=0.2 and TPR=0.7 indicates that an adversary can identify a truck with 70% success rate if it is indeed a truck. 20% of the time, the adversary will mis-identify an automobile as a truck. The 10-fold cross validation as shown in Fig. 9(b) and Fig. 9(c) evaluates the learning model by rotating training and testing data (in 9:1 ratio) 10 times. Although the Naive Bayes model does not perform very well, as can be inferred from Fig. 9(b), using QDA, an adversary can achieve higher success rate (about 96%) when he/she only focuses on tracking 40% of the trucks which show the most evidence of being a truck. This is a considerable improvement of the tracking ability of an adversary. For instance, in a group of ten vehicles containing one truck and nine automobiles passing through a mix-zone, an adversary only has about a 30% failure rate to mis-classify one of the nine automobiles as a truck while the real truck is able to be identified in 40% of cases. This dramatically reduces the protection of a mix-zone for such outlier vehicles, since the mix-zone model predicts a 90% failure rate. The QDA learning model can also help an adversary when there are multiple trucks in the anonymity group. For example, if there are two trucks in ten vehicles, an adversary only has a 28% failure rate to mis-classify one out of the eight automobiles as a truck while it has a 40% success rate to identify the real trucks. With an overall success rate of more than 30%, he/she can identify the right trace of the target truck.

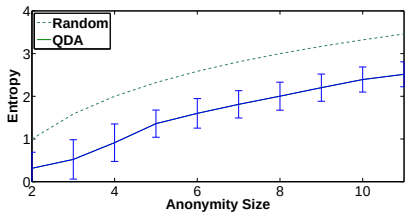

Figure 11: Entropy reduction due to outlier detection.

Fig. 9(d), Fig. 9(e) and Fig. 9(f) show the results when the headway information (which has been discussed in section 3.3.2) is available. Both training and testing results show that with headway information, an adversary will have a better chance to detect a truck under all three models. For example, as shown in Fig. 9(e), an adversary can achieve roughly 1% FPR under 35% TPR. This means that in a group of ten vehicles, one of which is a truck, passing through a mix-zone, the adversary only has a 8.6% chance to mis-classify one automobile as a truck while he/she is able to identify the real truck at a rate of 35%. This compares to a 90% failure rate with random guessing and shows that the protection of the mix-zone has deteriorated.

In Fig. 10, we assume that an adversary will perform a principal component analysis on the dataset. This operation projects all the characteristics (features) onto the 10 most important dimensions. The boundary is then generated only based on the 10 dimensions. Compared to Fig. 9, the performance is slightly better in both the training dataset and the testing dataset. This indicates that an adversary can achieve better results if he/she has a better learning model or data mining method. The better an adversary can do, the worse the current mix-zone model will be. In additional to all of the above observations, we also note that among all these three models, QDA performs the best and the headway information improves the performance of every learning model. We interpret this as evidence that further improvement is possible on outlier detection techniques if the adversary has better learning models or data mining algorithms.

In our study, we also evaluated the outlier detection techniques on motorcycle identification. Similar results were observed. However due to space limitations, the results are omitted from this paper.

Since entropy is one of the most important factors in evaluating

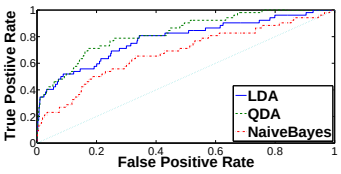

(a) performance comparison on the training dataset

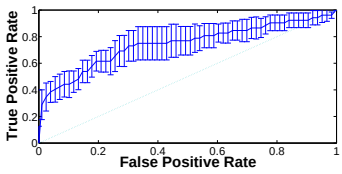

(b) 10-fold cross validation results for QDA learning model

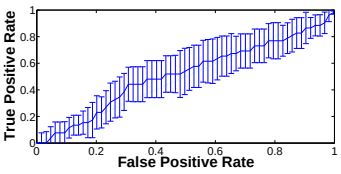

(c) 10-fold cross validation results for Naive Bayes learning model

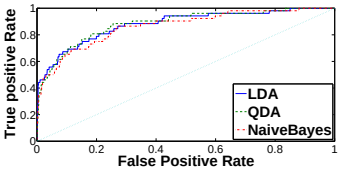

(d) performance comparison on the training dataset

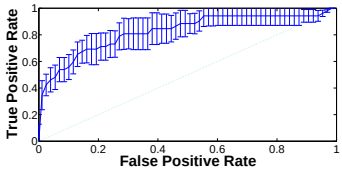

(e) 10-fold cross validation results for QDA learning model

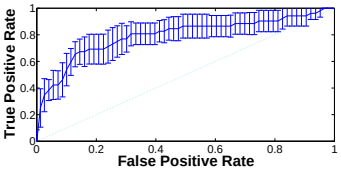

(f) 10-fold cross validation results for Naive Bayes learning model

Figure 9: Detecting truck from automobiles with manal feature selection

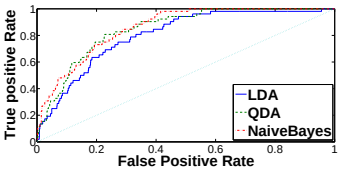

(a) performance comparison on the training dataset

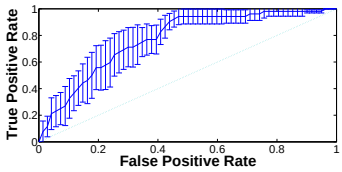

(b) 10-fold cross validation results for QDA learning model

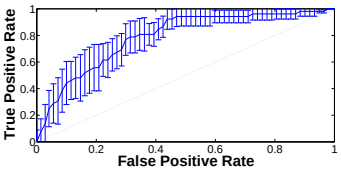

(c) 10-fold cross validation results for Naive Bayes learning model

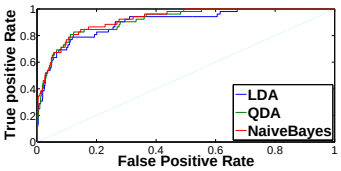

(d) performance comparison on the training dataset

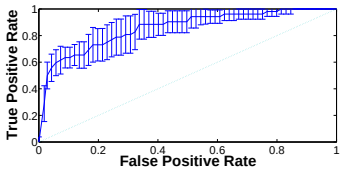

(e) 10-fold cross validation results for QDA learning model

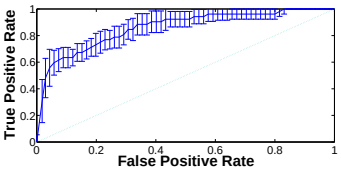

(f) 10-fold cross validation results for Naive Bayes learning model

Figure 10: Detecting trucks from automobiles using PCA projection on the first 10 dimensions.

a mix-zone model, in Fig. 11, we compare the original system entropy of a mix-zone with the one after the QDA learning model is used. As can be seen from the figure, while varying the size of the anonymity set, the QDA learning model always reduces the system entropy by at least 1. This reduction helps the adversary significantly, considering that entropy is a logarithmic value.

Finally, in another set of experiments, we create 55 groups of ten vehicles. Each group has one truck and the remaining are all automobiles which are moving closest to the truck. With the QDA model, an adversary can successfully identify and track trucks in the 22 of the 55 groups.

5.2 More General Outlier Detection: Extrinsic

Last, in this subsection, we show preliminary results on more general outlier detection: extrinsic and/or mixed cases. In this experiment, all 2000 feet long trace segments are further divided into three portions. The vehicle trace dataset from first and third portions are assumed to be available by the adversary. All trace dataset from the second portion of the road are removed in order to simulate a mix-zone model. Assume there is no intrinsic knowledge available such as vehicle type. The dataset from the first segment are used as training data. The adversary tries to link all the pairs of the traces belonging to the same vehicle. The learning model used is QDA. Different from previous work, in this part we assume the target can be any vehicle.

As shown in Fig. 12, the tracking rate, which is defined as the success rate at which an adversary can identify a particular target from a group of vehicles is largely increasing with the learning model compared to random guessing. For instance, in a ten vehicle anonymity set, the learning model has a 28% tracking rate, which is a good improvement for an adversary who originally had only a 10% tracking rate. As mentioned before, we define the anonymity set size as the number of nearby vehicles whose trace data are available to the adversary.

This figure has shown the tracking rate for any randomly picked

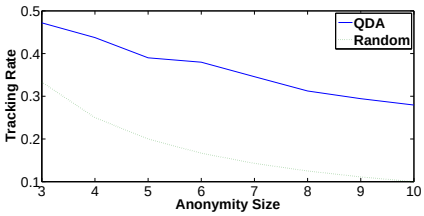

Figure 12: With the proposed generic outlier detection method, tracking rate is increased.

vehicle from the dataset. This indicates that there are a large number of cases in which the target may not have special movement characteristics compared to other vehicles. In practice, however, the adversary can ignore those cases and concentrate on outlier vehicles that can be tracked with high confidence.

6. CONCLUSION

In this paper, we have studied whether existing models to measure the degree of privacy in anonymized location traces hold as location traces continue to become more precise. Using data captured from vehicles we have shown that fine-grained location traces reveal speed distribution and acceleration patterns that can be used to distinguish traces from different vehicle types (e.g., trucks and cars). Our analysis on NGSIM location trace shows that an adversary can identify 40% trucks from cars with success rate of 96%. We have also shown that it is possible to identify outlier driving patterns such as higher speed, which could be used to link anonymous segments of location traces and eventually recover complete trips. Our preliminary results show that the general outlier detection technique can improve an adversary's ability to identify a trace segment of any user from an average tracking rate of 10% to 28%. While this rate is still relatively small, and would be smaller still if a vehicle trip passes over multiple mix zones, these findings show that movement characteristics reveal information. An immediate countermeasure is to revert back to coarser location traces but a full solution to this issue remains an open problem. We believe that further research is warranted to refine the definition of unlinkability for very fine-grained location traces.

7. REFERENCES

[1] National transportation statistics 2012. online, accessed Sept. 14, 2012.

[2] Next generation simulation (ngsim). online, accessed Sept. 14, 2012.

[3] H. X. Baik Hoh, Macro Gruteser and A. Alrabady. Achieving guaranteed anonymity in gps traces via uncertainty-aware path cloaking. In *IEEE Transactions on Mobile Computing 9 (8)*, 2010.

[4] H. X. Baik Hoh, Marco Gruteser and A. Alrabady. Enhancing security and privacy in traffic-monitoring systems. *Pervasive Computing, IEEE*, 5(4):38 –46, 2006.

[5] A. R. Beresford and F. Stajano. Mix zones: User privacy in location-aware services. In *In Proc. of the 2nd IEEE Annual Conference on Pervasive Computing and Communications Workshops (PERCOMW04)*, pages 127–131, 2004.

[6] L. Buttyan, T. Holczer, and I. Vajda. On the effectiveness of changing pseudonyms to provide location privacy in vanets. In *Proceedings of Workshop on Security and Privacy in Ad hoc and Sensor Networks*, 2007.

[7] D. Chaum, C. O. T. Acm, R. Rivest, and D. L. Chaum. Untraceable electronic mail, return addresses, and digital pseudonyms. *Communications of the ACM*, 24:84–88, 1981.

[8] T. Fawcett. Roc graphs: Notes and practical considerations for researchers. Technical report, HP Laboratories, 2004.

[9] J. Freudiger, M. Raya, M. FÍęlegyhÍćzi, P. Papadimitratos, and J.-P. Hubaux. Mix-Zones for Location Privacy in Vehicular Networks. In *ACM Workshop on Wireless Networking for Intelligent Transportation Systems (WiN-ITS)*, Vancouver, 2007.

[10] B. Gedik and L. Liu. Location privacy in mobile systems: A personalized anonymization model. In *Distributed Computing Systems, 2005. ICDCS 2005. Proceedings. 25th IEEE International Conference on*, pages 620 –629, 2005.

[11] P. Golle and K. Partridge. On the anonymity of home/work location pairs. In *Proceedings of the 7th International Conference on Pervasive Computing*, Pervasive '09, pages 390–397, Berlin, Heidelberg, 2009. Springer-Verlag.

[12] F. E. Grubbs. *Procedures for Detecting Outlying Observations in Samples*. Defense Technical Information Center, 1974.

[13] M. Gruteser and D. Grunwald. Anonymous usage of location-based services through spatial and temporal cloaking. In *ACM MobiSys*, pages 31–42, 2003.

[14] T. Hastie, R. Tibshirani, and J. Friedman. *The elements of statistical learning: data mining, inference and prediction*. Springer, 2 edition, 2009.

[15] B. Hoh and M. Gruteser. Preserving privacy in gps traces via uncertainty-aware path cloaking. In *Proceedings of ACM CCS*, 2007.

[16] B. Hull, V. Bychkovsky, Y. Zhang, K. Chen, M. Goraczko, A. K. Miu, E. Shih, H. Balakrishnan, and S. Madden. Cartel: A distributed mobile sensor computing system. In *4th ACM SenSys*, Boulder, CO, November 2006.

[17] J. Krumm. Inference attacks on location tracks. In *Proceedings of the Fifth International Conference on Pervasive Computing (Pervasive), volume 4480 of LNCS*, pages 127–143. Springer-Verlag, 2007.

[18] U. Lee, E. Magistretti, M. Gerla, P. Bellavista, and A. Corradi. Opportunistic dissemination and harvesting of urban monitoring information in vehicular sensor networks. In *UCLA, Tech. Rep*, 2007.

[19] U. Lee, B. Zhou, M. Gerla, E. Magistretti, P. Bellavista, and A. Corradi. Mobeyes: smart mobs for urban monitoring with a vehicular sensor network. *Wireless Communications, IEEE*, 13(5):52 –57, 2006.

[20] M. Li, K. Sampigethaya, L. Huang, and R. Poovendran. Swing & swap: user-centric approaches towards maximizing location privacy. In *Proceedings of the 5th ACM WPES' 06*.

[21] M. F. Mokbel, C. yin Chow, and W. G. Aref. The new casper: Query processing for location services without compromising privacy. In *In VLDB*, pages 763–774, 2006.

[22] T. Nadeem, S. Dashtinezhad, C. Liao, and L. Iftode. Trafficview: Traffic data dissemination using car-to-car communication. *ACM SIGMOBILE Mobile Computing and Communications Review*, 8, 2004.

[23] B. Zan, P. Hao, M. Gruteser, and X. Ban. Vtl zone-based path cloaking algorithms. In *Proceedings of the 2011 IEEE 14th International Conference on Intelligent Transportation Systems (ITSC 2011)*, 2011.