Learning Location from Shared Elevation Profiles in Fitness Apps: A Privacy Perspective

Ulku Meteriz-Yildiran, Necip Fazil Yildiran, Joongheon Kim, and David Mohaisen

Abstract—The extensive use of smartphones and wearable devices has facilitated many useful applications. For example, with Global Positioning System (GPS)-equipped smart and wearable devices, many applications can gather, process, and share rich metadata, such as geolocation, trajectories, elevation, and time. For example, fitness applications, such as Runkeeper and Strava, utilize the information for activity tracking and have recently witnessed a boom in popularity. Those fitness tracker applications have their own web platforms and allow users to share activities on such platforms or even with other social network platforms. To preserve the privacy of users while allowing sharing, several of those platforms may allow users to disclose partial information, such as the elevation profile for an activity, which supposedly would not leak the location of the users. In this work, and as a cautionary tale, we create a proof of concept where we examine the extent to which elevation profiles can be used to predict the location of users. To tackle this problem, we devise three plausible threat settings under which the city or borough of the targets can be predicted. Those threat settings define the amount of information available to the adversary to launch the prediction attacks. Establishing that simple features of elevation profiles, e.g., spectral features, are insufficient, we devise both natural language processing (NLP)-inspired text-like representation and computer vision-inspired image-like representation of elevation profiles, and we convert the problem at hand into text and image classification problem. We use both traditional machine learning- and deep learning-based techniques and achieve a prediction success rate ranging from 59.59% to 99.80%. The findings are alarming, highlighting that sharing elevation information may have significant location privacy risks.

Index Terms—location privacy, privacy breach, privacy in social media, fitness applications, natural language processing, applied machine learning

✦

1 INTRODUCTION

From smartphones to wearables, an increasing number of Internet of Things (IoT) devices are equipped with Global Positioning System (GPS), accelerometers, and gyroscopes to allow applications to function or to present a better user experience using *geodata*, such as location and elevation information. More recently, fitness applications that run on smartphones and smartwatches used these components to collect spatial, temporal, and activity-specific information to analyze, summarize, and visualize users' activities. By analyzing each activity, many of those applications deliver personalized motivations and challenges for users to meet their goals. Using social media support of these applications for sharing updates about users' activities, including training routes and elevation profiles for the routes taken for an activity (e.g., walking, running, climbing, cycling), users can have positive behavioral changes through a more active lifestyle motivated by competitions with acquaintances [2].

Despite the broad set of advantages that geodata offers, geodata usage and uncontrolled sharing can pose a signif-

- *U. Meteriz-Yildiran is with Meta and N. Yildiran is with Google; the work of both authors was done while they were at the University of Central Florida. D. Mohaisen is with the Department of Computer Science, University of Central Florida, Orlando, FL 32816, USA. J. Kim is with the Department of Electrical Engineering at Korea University, Republic of Korea. D. Mohaisen is the corresponding author (e-mail: mohaisen@ucf.edu). An earlier version of this work has appeared in IEEE ICDCS 2020 [1]. This work was supported in part by NRF under grant 2016K1A1A2912757 and CyberFlorida Seed Grant (2021/2022). J. Kim was supported by NRF under grant 2022R1A2C20048690. J. Kim (joongheon@korea.ac.kr) and D. Mohaisen (mohaisen@ucf.edu) are the corresponding authors.*

Fig. 1. Survey results for understanding users' behavior with starting point statistics and finishing point statistics. While 90% of the 60 participants indicated their start of activity is either home, school, or work, an overwhelming 98% of the participant indicated those to be the finish (end) point of their activities.

icant privacy risk that can be further exploited in multiple attacks, including stalking [3] and cybercasing [4]. For example, with a large amount of geotagged data, including text, images, and videos, cybercasing provides criminals and maliciously motivated individuals with a significant attack vector. Geo-tagged photos that are frequently posted on image-sharing websites, such as Flickr, or second-hand sale websites, such as Craigslist, may put owners of those images at risk. For example, geo-tagged images posted on sales websites may reveal the location of the advertised product, leading to trespassing or even theft.

While geodata recorded by fitness applications is indeed important and valuable for the operation of those applications, this data can also be used for launching attacks on users by breaching their privacy since sensitive information of users, such as home or workplace location, can be easily inferred from such data. Even worse, a large number of users, when sharing such information, would be unaware of the ramifications of sharing and the potential risk of inferring such contextual information, such as home, work

location, etc., from such shared location data. To support this argument, we conducted an online survey with 60 participants who regularly use fitness applications outdoors. The results of the survey, summarized in Figure 1, reveal that 51% of the participants start their training from their homes, 36% start from their school, and 3% start from their workplace, while 76% of the participants finish their training at their homes. Moreover, for the same set of users (results are not shown in Figure 1), 42% of those users have indicated that not sharing location information implies privacy protection, while 30% of the respondent were uncertain, and 28% were certain that not sharing would not necessarily mean their privacy is protected. The mixed responses highlight the gap between reality and expectations of privacy when sharing location information online and call for further investigation.

Although it is possible to hide the location trajectory by removing the activity map in the fitness applications, users still want to share elevation profiles or certain statistics of the activity to show the roughness, technicality, and difficulty of the routes they took as a measure of their workout. For example, up until recently, users have been demanding those fitness applications to allow for fine-grained and customized access control by allowing them to share the elevation profile of an activity while masking the map that highlights the actual trajectory, which is deemed of high privacy value to them [5]–[8].

In the same survey we conducted earlier, we asked our 60 subjects "while sharing an outdoor workout record, do you think hiding the map and sharing only the statistics of your training (such as speed and elevation changes) is enough for protecting your privacy?". The results were overwhelmingly positive, with 25 of them indicating "yes", 18 indicating "maybe" (together accounting for more than 71%), and only 17 indicating "no".

Is sharing the elevation profile of activity enough to maintain the privacy of users? In this paper, we argue that an approximate location, extracted from the contexts of activities and at different levels of location granularity, could still be revealed from the elevation profile information. We examine this problem comprehensively and develop techniques that can be used to accurately associate an elevation profile with contextual information, such as the location.

Contributions. In this paper, we contribute the following:

- we translate the problem of location privacy inference from elevation profiles into text classification and image classification problems by encoding the elevation signals as strings and visualizing the elevation signals as images to employ various common approaches for solving image and text classification problems,
- we investigate the possible attack surface for the problem by exploring three different threat models, which we later use to evaluate the success of our approaches by simulating our methods considering each threat model,
- we demonstrate that location information can be predicted from elevation profile using different machine/deep learning methods with accuracy in the range $80.25\% - 99.80\%$ at different resolutions.

We note that examining the effect of the attack using a large-scale in-the-wild case study is impractical as service

TABLE 1
Popular fitness applications and their features. ET: Exercise tracking. SS: Ability to share to social media. SNS: Social networking capabilities in the service. PR: Private records. BU: User blocking capability.

Service	ET	SS	SNS	PR	BU
Strava	●	●	●	●	●
Runtastic	●	●	●	○	●
Runkeeper	●	●	●	●	○
Nike+ Running	●	●	●	●	○
MapMyRun	●	●	●	●	○

providers prevent the use of their data for tracking by a third party. However, to motivate the effect of the attack, we consider the scenario of an informed adversary who knows the city where a victim with the exposed elevation profiles for associated activities lives. As such, the adversary proceeds by profiling the city and collecting elevation profiles for different segments within the city. One can see how easily such an adversary will be able to contextualize the elevation profiles of the victim's activity further by narrowing it down to a few candidate precomputed elevation profiles. Given the adversary's awareness of the mapping between the location and the profiles, the adversary will be able to easily infer valuable information about the habits of the victim by associating, for instance, end, start, and stopping points on the elevation profile, with points of interest (cafes, workplace, etc.).

Organization. We present the background in section 2, the threat model in section 3, a high-level overview of our approach in section 4, the implementation details are presented in section 5, the evaluation results in section 6, further discussions in section 7, the related work in section 8, and concluding remarks in section 9.

2 BACKGROUND

In this section, we provide some background information highlighting the significance of elevation profiles for athletes, the use cases, some properties of the fitness applications on the market today, and some reported privacy breach incidences of fitness applications to contextualize further the work presented in the rest of this paper.

2.1 Elevation Profiles Importance for Athletes

Athletes who keep track of their activity records measure various modalities and attributes associated with the activities, including the distance, speed, overall time, and heart rate over the course of the activity. Based on these attributes, they adjust their training strategies to reach their goals. Elevation changes, often reported in the form of elevation gain, are one of the most significant attributes measuring the performance of a cyclist/runner and often depict how hard the run or ride is. For example, riding a bike for a 20-mile ride while climbing 1000 feet in total is significantly more challenging than biking on a flat terrain [9]. Therefore, when recording or sharing a ride/run, athletes care about the changes in the elevation, thus elevation profiles.

2.2 Fitness Applications & Privacy Breach Incidents

Fitness applications allow users to track their workout history and provide them with statistics. Moreover, some fitness applications have social network capabilities, as shown in Table 1, and allow users to share workout summaries that are known to motivate users and their social network connections to achieve their goals [2]. Some fitness applications also inherit user-blocking features and capabilities from social network platforms, including user privacy options such as private records–the activity records that are only visible to the user.

Although fitness applications have configurable privacy options, there have been a lot of privacy incidents concerning location data obtained from those fitness applications. We review some of those privacy breaches in the following to contextualize our work in the broader privacy literature.

Revealing Secret U.S. Military Bases. Strava, which is one the most popular fitness tracking applications in the market today, collects users' public data and publishes a heatmap of the aggregates to highlight routes frequented by users [10]. Although the aggregates in the heatmap do not explicitly contain any identity information, activities in desolate places revealed the location of many U.S. military bases, which is considered sensitive information [11], [12].

Deanonymization Through Strava Segments. In Strava, the heatmap feature was used to show "heat" made by the aggregated and public activities of Strava users over the past year. It is, however, shown that a dedicated adversary can deanonymize heatmap to find out users who ran in a specified route [13]. For example, by selecting a route from the heatmap, a registered user can manually create a GPS eXchange (GPX) track file and create a segment using it on Strava. A segment is a portion of a road or a trail where athletes compare their finishing times. Consequently, once this segment is created, the users who previously ran that route are shown on the leaderboard grouped by gender and age. This feature is then leveraged to identify individuals who ran that particular place.

Tracking and Bicycle Theft. Users of fitness applications can share information related to the equipment used for the activity, including bicycles, tracking devices, shoes, etc., along with the routes frequented. The combined shared information makes them a target for robbery, and several such incidents of bicycle theft are reported [14]–[17].

Attack on Privacy Zone. To cope with the increasing privacy risks, Strava features *privacy zones*, a technique to obfuscate the exact start and end points of a route. A recent study [18] has demonstrated that it is possible to reveal the exact start and end point of a route that utilizes the privacy zone feature. The same study also claimed that around 95% of the users are at risk of revealing their location information.

Live Activity Breach. In Runtastic, one of the popular activity-tracking applications, users can share their live activities. In theory, users should be able to configure the privacy settings for their activities such that only privileged users, such as connections on the application platform, can track the shared live activity session. However, it has been demonstrated [19] that the selected privacy settings are not correctly applied to a live session. As a result, everyone can go through live sessions and track Runtastic users in real time, even though the associated privacy options should have prevented this type of breach. Based on this incident, it would be easy to stalk and locate a user, e.g., a lone runner or cyclist with expensive equipment, in real time.

3 THREAT MODELS

We outline the potential threat models under which this study is conducted. We describe three models under which location privacy is breached only from associated elevation profiles. We note that the following threat models are only hypothetical: no attacks were actually launched on any users. As mentioned earlier, this study in its entirety is motivated by the aforementioned demands of users to have more flexibility over-sharing partial data, such as elevation profiles, and examines the ramifications of such sharing in a hypothetical setting. We note, however, that those settings are also plausible if such sharing is enabled.

Our study utilizes three threat models: TM-1, TM-2, and TM-3, which we outline below with their justifications. The adversarial capabilities in TM-1 are greater than in TM-2 and TM-3, making it a more restrictive (powerful) model.

❶ TM-1. In TM-1, we assume an adversary with workout history records of a target user, and the goal of the adversary is to identify the last workout location of the target user from the recently shared elevation profiles. TM-1 is justified by multiple plausible scenarios in practice. For example, such an adversary might have been a previous social network connection of the target user that was later blocked. In such a scenario, the adversary may have previous workout records of the target from which the adversary may attempt to de-anonymize the target's activities. Another example might include group activities, where two individuals (i.e., the adversary and target) may have shared the same route at some point. In either case, by knowing the target's previous fitness activity records, the main goal of the adversary in this model is to identify recent whereabouts only from publicly shared elevation profiles in workout summaries, thus breaching the target's location privacy.

❷ TM-2. In TM-2, we assume an adversary with access to limited information, such as the city where the target lives. Such information is easily accessible from public profile summaries, athlinks.com, public records, etc. The adversary's goal in TM-2 is to find out which region or part of a given city the target's activities are associated with. The TM-2 use scenario may include a targeted user sharing private activities in which the route is hidden while the elevation profile is shown. The adversary, knowing the city where the target lives, would want to identify the region (e.g., a borough in the city) associated with the user's activity.

❸ TM-3. In TM-3, we assume an adversary trying to identify the target user's city using only publicly shared elevation profiles without any prior information. We assume, however, the adversary has the ability to profile the elevation of cities with information that is easily obtained from public sources (e.g., Google Maps, OpenStreetMap). The use scenario of TM-3 may be used as a stepping stone towards launching the attack scenario in TM-2 upon narrowing down the search space to a city.

4 APPROACH: HIGH-LEVEL OVERVIEW

In this section, we give a brief overview of our pipeline, which consists of the data collection, preprocessing, feature extraction, and classification as illustrated in Figure 2. Each phase of the pipeline is detailed in section 5.

Data Collection. We collected three datasets with varying and rich characteristics, namely (i) user-specific activity data collected from an athlete, (ii) mined training route segments grouped at city-level, and (iii) mined training route segments grouped at borough-level. For the user-specific dataset, we collected physical activity records of athletes and converted those activities to an intermediate format, the GPS Exchange Format (GPX). Then, we parsed the GPX files and manually labeled them according to the latitude and longitude information included within each file. For the second dataset, we mined training route segments from a popular fitness tracking website by specifying the location boundaries, *i.e.,* the class label of the mined data, and augmented each segment with the corresponding elevation profiles obtained from Google Maps Elevation API. Finally, we similarly constructed the borough-level dataset as in the city-level dataset.

Preprocessing. We employ Natural Language Processing (NLP) and computer vision techniques to convert the problem to text classification and image classification problems, respectively. To this end, we prepare the data accordingly in the preprocessing phase. Preprocessing consists of two parts: (i) text-like and (ii) image-like representations.

For text-like representation, we discretize the elevation signals and compute the minimum required *word* size. We then create a mapping between each unique discrete value and a string. By mapping the string correspondents to the unique discrete values, we encode the elevation profiles in text. We, then, form a *vocabulary* from the text sequences of each dataset using the n-grams.

To obtain image-like representations, we convert the elevation profiles to a fixed-sized line graph where the *x*-axis stands for time and the *y*-axis stands for the elevation values. We also color the lines in the graphs to represent the elevation interval in which the elevation profiles range.

Feature Extraction. The classification algorithms operate on high-quality and discriminative features obtained from the representations of elevation profiles. For feature extraction, we utilize NLP and computer vision approaches.

To employ NLP approaches using the vocabulary obtained in preprocessing phase, we represent each elevation profile as either a feature vector based on the vocabulary frequency in the text-like representation (bag-of-words vector) or as a term frequency-inverse document frequency (tf-idf) vector. To employ computer vision approaches, we utilize Convolutional Neural Networks (CNN) over image-like representations. The optimal features of an image-like representation are efficiently extracted by the convolutional and pooling layers in the CNN architecture.

Multi-Class Classification. We utilize various machine learning and deep learning models for classification, including Support Vector Machine (SVM) and Random Forest Classification (RF), Multi-Layer Perceptron (MLP), Long Short-Term Memory (LSTM), 1D Convolutional Neural Network (C1D), and 2D Convolutional Neural Network (CNN).

5 IMPLEMENTATION DETAILS

The implementation details of data collection, preprocessing, feature extraction, and multi-class classification are addressed in the following subsections.

5.1 Data Collection

In this study, we compiled three datasets: the user-specific dataset, the city-level dataset, and the borough-level dataset. The user-specific dataset is retrieved from a voluntary athlete who frequently records activities through fitness applications. It offers dense and thorough coverage of regions frequented by the user; those regions are used as class labels. The city-level and borough-level datasets are created from scratch by collecting location trajectories that are created and frequented by the athletes. Both city-level and borough-level datasets provide sparse coverage of cities and boroughs.

5.1.1 User-Specific Dataset

For the user-specific dataset, we collected activity data, including each activity's location trajectory and the corresponding elevation profile from a voluntary athlete who records activities frequently through fitness applications. First, the location trajectories included in the user-specific dataset are converted to GPX format to avoid confusion caused by different formats and settings across the activity records. Then, to label the samples, the maximum and minimum coordinates of each location trajectory are fetched. Each sample location trajectory is encapsulated with a tight rectangle whose top right (North East) and bottom left (South West) corners are computed from the maximum and minimum coordinates of the trajectory as illustrated in Figure 4. To classify the samples, each rectangle encapsulating the trajectory is compared with the previously created regions. If the Euclidean distance between the center of the rectangle and the center of the existing region does not exceed a predetermined threshold, the rectangle and its corresponding sample are labeled with a unique identity of the region. Then, we annotated the region labels, such as Orlando, Washington DC etc., based on the manual observation on the map. If no region includes the trajectory, a new region is created. The final sample size distribution of the user-specific dataset is shown in Table 2.

The user-specific dataset is prone to have similar location trajectory portions across its samples since the user may frequent the same set of places in his/her everyday activities, such as the location trace they follow while leaving/arriving home or their favorite routes. Therefore, we calculated the average overlap ratio of the routes included in the user-specific dataset by comparing each sample with the other samples with the same class label. For each sample pair comparison, the overlap ratio is calculated as the intersection-over-union of the tight rectangles encapsulating the sample routes. The average overlap ratio of the user-specific dataset is calculated as 35%.

5.1.2 City-Level Dataset

For the city-level dataset, we mined **publicly available** training route segments in a popular fitness tracking application using its EXPLORESEGMENTS() functionality. We note that our experiments do not put any users at risk and
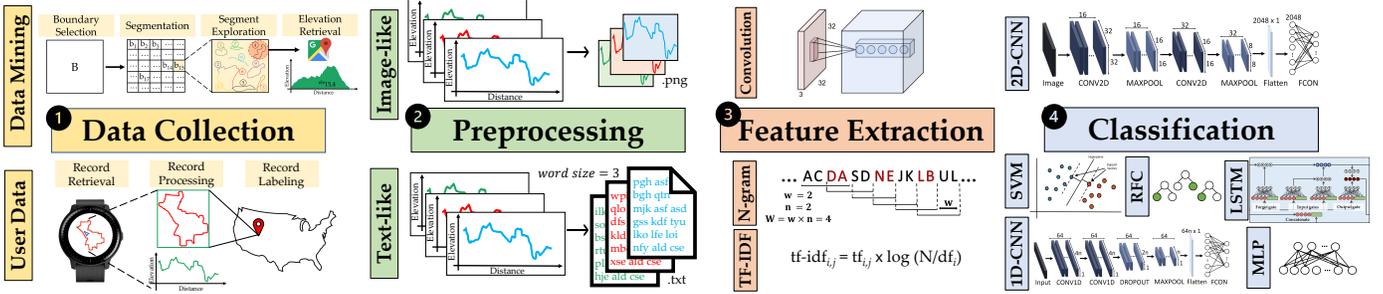
Fig. 2. The end-to-end pipeline of the approach consists of four main stages: (i) data collection, (ii) preprocessing, (iii) feature extraction, and (iv) classification. There are two types of data representations, each of which is processed differently in the feature extraction and classification stages. The feature extraction of the image-like representation is internally handled in the convolution layers of the classification phase. The convolution illustration in the feature extraction step is for the sake of modularization and for consistency with the other pipeline instance.
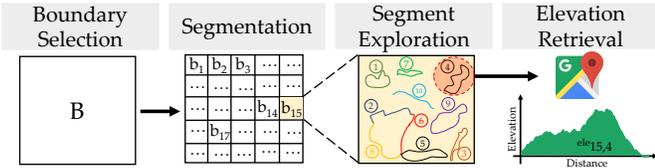
Fig. 3. An illustration of the data mining pipeline. A geolocation boundary, B, is segmented into small boundaries, each of which is then forwarded to the Segment Exploration step to obtain the most frequented ten route segments for that particular boundary. Finally, the elevation profile of each route segment is retrieved from the Google Maps API.

Fig. 4. An illustration of the tight rectangle encapsulating an example route. The rectangle is created by fitting the route in between the minimum and maximum $(latitude, longitude)$ pairs of the given route. The minimum and maximum $(latitude, longitude)$ pairs correspond to the bottom left (*i.e.,* South West) corner and the top right (*i.e.,* North East) corner of the rectangle, respectively.

are not in violation of the terms of use of the fitness tracking application: since both the trajectory (map) and elevation profiles are public information, we are also not obtaining any information beyond what is provided by the users explicitly. We also note that the training route segments are user-created activity routes whose main purpose is to compare completion times among users who also completed the same route. They are particularly useful for our purposes since they include public location trajectories that are frequented by the actual users rather than randomly created location trajectories that may not necessarily be of privacy value. During the segments mining, the anonymity—thus the privacy—of the users who frequented the segments or created the segments is maintained.

The overall data mining procedure consists of three steps, as illustrated in Figure 3. First, we define the cities of interest, which we also use as the class labels per our threat model. For each city, we define the rectangle geolocation boundary box B consisting of the top right and bottom left corner coordinates in the boundary selection phase. In the segmentation phase, and since EXPLORESEGMENTS() returns only the ten most frequented segments encapsulated by a given boundary, and to obtain more data of a geolocation boundary box, we divide the large rectangle boundary of the city into smaller region boundaries, each denoted by b_i, by following a grid-like structure as shown in the second phase of the Figure 3. For each region boundary b_i, we call EXPLORESEGMENTS() and receive the geolocation polyline path, $path_i^j$ where $j \in [1, 10]$, of the 10 most frequented segments encapsulated in b_i, as shown in the segment exploration phase. Finally, since the polyline paths do not include elevation profiles, we

TABLE 2
User-specific dataset sample size distribution.

Regions	Abbreviation	Sample Size
Washington DC	WDC	366
Orlando	ORL	232
New York City	NYC	120
San Diego	SD	18

obtain the associated elevation profile $ele_{i,j}$ for each $path_i^j$ using the Google Maps Elevation API through the elevation retrieval phase. The sample size distribution of the city-level dataset can be found in Table 3.

Unlike the user-specific dataset, the city-level dataset does not include any overlapped samples since each region r_i is disjoint with the other regions. A segment route included by more than one neighboring region is not considered since EXPLORESEGMENTS() returns the routes encapsulated within the given boundaries, b_i.

5.1.3 Borough-Level Dataset

For the borough-level dataset, we apply a similar mining procedure as we have done with the city-level dataset, using the borough boundaries instead of the city boundaries. Table 4 shows the sample size distribution of the borough-level dataset for different cities.

TABLE 3
City-level dataset sample size distribution.

Regions	Abbreviation	Sample Size
New York City	NYC	2437
Washington DC	WDC	2129
San Francisco	SF	743
Colorado Springs	CS	369
Minneapolis	MIN	363
Los Angeles	LA	280
New Jersey	NJ	266
Duluth	DUL	156
Miami	MIA	94
Tampa	TAM	83

TABLE 4
Borough-level dataset sample size distribution.

City/State	Region	Sample Size
Los Angeles (LA)	Downtown	280
	Santa Monica	128
	Chinatown	46
	Beverly Hills	38
Miami (MIA)	Downtown	67
	Miami Beach	44
	Virginia Key	18
New Jersey (NJ)	Jersey City	266
	West New York	23
	Newark	28
New York City (NYC)	Manhattan	2437
	Queens	353
	Brooklyn (South)	239
	Brooklyn (North)	205
	Bronx	142
	Staten Island	119
San Francisco (SF)	South West	743
	South East	144
	North West	130
	North East	86
Washington DC (WDC)	District of Columbia	2129
	Baltimore	218

5.2 Preprocessing

A key design element in our pipeline is the representation modality of the elevation profile, which will significantly impact the performance of our elevation-location mapping, as we show later. We transform the samples into text-like and image-like representations to facilitate feature extraction and feed them into our classification module. In this section, we describe the details of the utilized preprocessing methods.

5.2.1 Text-like Representation

For our text-like representation, our approach follows four steps, as shown in Figure 5: discretization, word size decision, text encoding, and vocabulary creation.

Discretization. In the discretization step, the original elevation signal is discretized along the y-axis, which represents the elevation values to avoid possible overhead by small differences in the precision causing longer string correspondences and, consequently, longer vocabulary and sparse feature vectors. The discretization is done as follows. Let e_i^j be the i-th elevation value in j-th sample. The discretization functions are defined as $f(e_i^j) = \lfloor e_i^j \rfloor$ and $f(e_i^j) = \frac{\lfloor e_i^j \times 10^3 \rfloor}{10^3}$, where the first function is used for processing the user-specific dataset and the second function is used for processing the city-level and borough-level datasets. Since the user-specific dataset is dense in terms of sampling rate, using the floor function is enough to represent the routes. However, as the city-level and borough-level datasets are already sparse, losing information is undesired, so we used the second function to represent the elevations that differ in up to 3 decimal digits precision. Having more fine-level details when the samples are sparse would help our models to learn more and thus increase the accuracy. To demonstrate the effect of discretization, we measured the vocabulary size of the smallest class of the User-specific dataset, i.e., San Diego. For a class as small as the San Diego class, using the second function results in a vocabulary of size 12,870, and using the first function results in a vocabulary of size 3,155. Such difference in the vocabulary size demonstrates the effect and necessity of discretization.

Word size decision. For the word size decision, we use $w = \log_l c$, where w is the word size, l is the length of the alphabet, and c is the number of unique value occurrences in the given signals.

Text encoding. For text encoding, each unique value in all the discrete signals is mapped to a unique string with length w, and each sample signal is encoded by referring to the string correspondences of each value in the discrete signal and concatenating these strings to construct a long text, *i.e.,* corpus.

Vocabulary creation. To create our *vocabulary*, we consider the corpus created from all encoded signals regardless of labels. Each line in the corpus represents a sample signal, and each word in a line represents the text correspondence of an elevation value in the sample signal. We build a vocabulary from the unique word-based n-grams of the document. As illustrated in Figure 6, a window with size $W = w \times n$ is slid throughout the corpus and each window content is appended to the vocabulary set. Since the vocabulary set does not contain duplicate entries by definition, we construct the vocabulary consisting of unique n-grams of the given dataset after traversing the corpus by n times with different window sizes.

5.2.2 Image-like Representation

In the image-like representation, the elevation signals are drawn as line graphs and saved as a 32×32 images[1]. To draw a line graph, the maximum and minimum values for the y-axis are set to be the maximum and minimum of each elevation signal, and the lines are colored to encode the value interval in which the elevation signal ranges. We note that the image-like representation with colors has multiple

---

1. The size is chosen to strike a balance between the performance in terms of the required computations and produced accuracy.
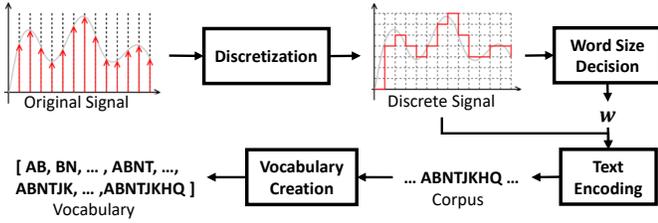
Fig. 5. Illustration of the flow of text-like preprocessing. The signal is discretized by eliminating the small elevation fluctuations. The discretized signal is also used for deciding the word size of the encoding. The discrete signal is then encoded in text and a vocabulary is built.
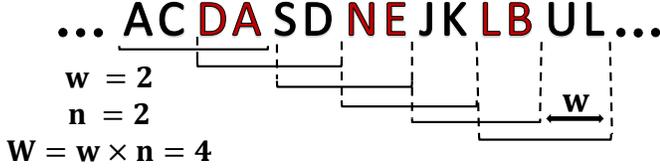
Fig. 6. Illustration of bi-gram creation where the word size is $w = 2$ and window size is $W = 4$.

(a) New York City B/W
(b) Los Angeles B/W
(c) New York City colored
(d) Los Angeles colored

Fig. 7. *Elevation profile graphs by fixing the y-axis range and using only black versus elevation profile graphs by fitting and using color encoding.* As can be seen, with the black option the elevation range can be represented by the position of the signal, but changes in the signal are not visible, which is an information loss. However, with the color option enabled, both alterations in the signal and the signal range are represented in one image.

advantages over the black-white representation where the y-axis is set to the range of a whole dataset. First, as illustrated in Figure 7, the alterations of an elevation signal are more visible with the color encoding method, which could be a more discriminative feature to learn. Second, the color encoding method results in efficient utilization of the feature space. We use 200 elevation values for each image, obtained by dividing the elevation signal into equal-sized parts.

5.3 Feature Extraction

To classify elevation profiles accurately, we extract discriminative features from the elevation profile representations.

Text-like. In the text-like feature extraction, we utilized two methods: (i) n-grams, (ii) tf-idf.

For the n-grams method, words and non-overlapping occurrences of word sequences are counted, and a feature vector for each sample is created with each unique word sequence count being a feature. Finally, the feature vectors are normalized where each feature represents the occurrence probability of each word in the given sample.

The tf-idf is a statistical feature signifying the importance of a word in a document. The tf-idf values proportionally increase as the number of appearances of a word in a document increases. Technically, the tf-idf for a word is the multiplication of two metrics: (i) term frequency (tf) and (ii) inverse document frequency (idf). Tf-idf of a word t in a document d included in the set of documents D of which cardinality is N is calculated as follows:

$$\text{tf-idf}(t, d, D) = \text{tf}(t, d) \cdot \text{idf}(t, D),$$
$$\text{tf}(t, d) = \log(1 + \text{freq}(t, d)),$$
$$\text{idf}(t, D) = \log\left(\frac{N}{\text{count}(d \in D : t \in d)}\right).$$

The higher the tf-idf means that the word t is more relevant to the particular document d.

Feature Selection. When the dataset is large and diverse, the vocabulary and, consequently, the feature vector representation become too large to process, compute, and learn from. With a feature selection phase to address long feature vectors, some rarely occurring features in the vocabulary are discarded according to a pre-specified feature frequency threshold. For selection, the features are ordered by their term frequency across the corpus, the features whose term frequency is below a specified threshold are discarded, and a new vocabulary is created. For both feature extraction methods of the text-like representations, the term frequency threshold is set such that the size of the eventual vector representation is 5,000.

Image-like. We use the CNN for processing and classifying the 32×32 images, thus it is unnecessary to explicitly extract features since the convolutional layer kernels do that already by learning the filters optimally and efficiently. Therefore, the actual feature extraction mechanism for the image-like representation is discussed in the context of the classification phase.

5.4 Multi-Class Classification

For classification, SVM, RF, MLP, LSTM, and CNN are used.

SVM. SVM is a supervised classification technique. The main challenge in SVM is finding the best hyperplane that divides the classes from each other considering a given margin. SVM with the linear kernel is able to distinguish the classes more successfully when the features are multidimensional and numerous [20], [21]. In linear kernel settings, when an optimal hyperplane is found while representing a class, only the features around the hyperplane within the given margin are considered and the other features are simply ignored. As such, the complexity of SVM is

7
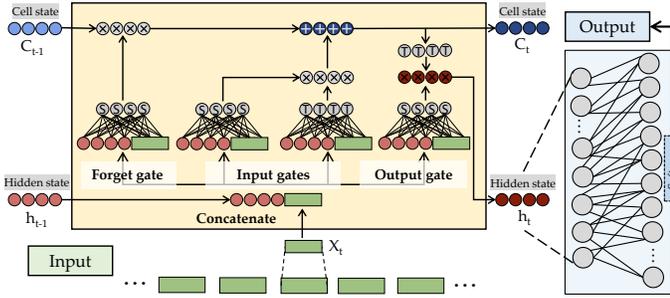
Fig. 8. The architecture of the LSTM network consists of an LSTM unit with four hidden layers and two fully-connected layers. The input samples are passed through the LSTM unit individually and the last hidden state of the LSTM unit is forwarded to the fully-connected neural network. The fully-connected neural network utilizes softmax activation at the output layer which outputs the class probabilities.

(a) 1-D CNN Architecture

(b) 2-D CNN Architecture

Fig. 9. The CNN architectures used for classification. A 1-D CNN is used for the one-dimensional data representations, i.e., n-grams, tf-idf, and raw data. A 2-D CNN is used for the image-like representation. The numbers in the figures depict the dimensions of the input throughout the learning/prediction process.

independent of the number of features. Since we consider n-grams up to $n=5$ in the feature extraction phase, the number of features is numerous which makes the usage of SVM legitimate in terms of efficiency and success.

SVM with a linear kernel is generally used for binary classification. To utilize it for the multi-class classification problem, we use the one-versus-rest method. In this method, an individual model is trained for each class and the label of the most confident model is outputted. For the penalization, we use the L2 norm, as it is a standard for linear SVMs. As for the loss function, we utilized the square of hinge loss, and the hyperparameters are decided through grid search tuning.

RF. RF is based on decision trees, which are ensemble learning methods for classification. The main feature of ensemble methods is further improving the generality and robustness of a single estimator by combining several base estimators that are built with a given learning algorithm. In decision trees, features are represented by tree nodes and each branch between two nodes represents what the immediate ancestor node returned. Since building an optimal binary decision tree from given features is an NP-complete problem, using a Random Decision Forest with different tree configurations and efficient heuristics is a way to alleviate the NP-completeness for classification problems. While creating our random forest, perturb-and-combine techniques are used. Perturb-and-combine techniques are designed specifically for decision trees to improve their accuracy by creating several (different) versions of the estimator by *perturbing* the training set, then *combining* these different versions into a single estimator [22]. Further details on this approach can be found in [23] and [22].

In this study, we use 20 decision trees for the RF model. The final prediction is then done by averaging the tree predictions. We do not set any upper limit on the number of leaf nodes or the depth of the tree, i.e., there were no time optimization concerns during training the process, so we leave the trees to grow to their maximum depth.

MLP. MLP is a feed-forward fully-connected neural network that utilizes backpropagation for training and is used for supervised learning.

Recent studies on comparing multi-layer neural networks and decision trees [24], [25] concluded the following:
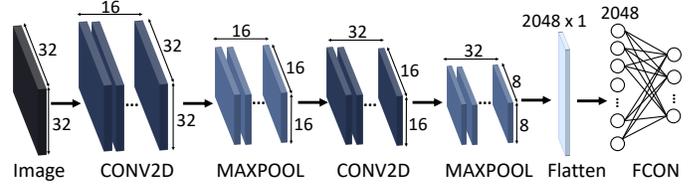
- Multi-layer neural networks allow incremental learning, in which the model's knowledge is continuously extended, more easily than the decision trees.
- The training time of the multi-layer neural networks is much longer than decision trees.
- The multi-layer neural network predictions are generally as good as the predictions produced by the decision trees, although they can perform better in certain cases.

Given the aforementioned potential for MLP to perform better than RF, we used MLP with 20 hidden layers in our experiments. We used Adam solver [26] for weight optimization since it is shown to perform well in terms of both training time and validation score for large feature spaces. We also utilize ReLU as the activation function, 0.001 as the learning rate, and 200 as the epoch size. For regularization, we use L2 norm with 0.0001 penalty parameter. The hyperparameters are decided through grid search tuning.

LSTM. LSTM is a recurrent neural network architecture. Unlike the standard neural networks, LSTMs are capable of keeping track of long-term dependencies in the input sequences using feedback connections. Such long dependencies are handled through feature extraction in n-grams and tf-idf vectors for standard neural networks. LSTM, on the other hand, handles dependencies internally, without requiring an explicit feature extraction. LSTM is also particularly useful for capturing the order dependence in sequence prediction problems. For LSTM, the input sequence can be a time series, a sentence from a given language, or a text-like representation as in our application case. Figure 8 depicts the LSTM architecture employed in this work, where the input is directly passed through the LSTM layer with four hidden layers, and two fully-connected layers following the LSTM layer. Each input sample (vector representation of an elevation profile for n-grams and tf-idf, or individual values for the raw data) is passed through the LSTM unit. When all sample vectors (or elevation values) are passed through the LSTM units, the final hidden state vector is

passed to the subsequent fully-connected layer. In the LSTM unit, hyperbolic tangent and sigmoid (depicted as T and S, respectively, in Figure 8) are used as the activation function. For the fully-connected layers, ReLU and softmax activation functions are used, respectively.

CNN. CNN is similar to neural networks in the mechanism, as both of them consist of neurons with learned parameters, weights, and biases. The improvement of CNN, however, is in the form of convolution layers, which apply forward passes that decrease the number of parameters of the neural network considerably. Convolution layers also facilitate the processing of high-dimensional data, such as images. They prepare high-dimensional data for a fully-connected layer, which cannot process high-dimensional data efficiently, by highlighting the important spatial features along the way.

In this study, we utilize two CNN architectures, differing in the convolution layers dimensions. Figure 9 illustrates the employed CNN architectures.

In the first architecture, we use two consecutive 2D convolution layers (CONV2D) along with the ReLU activation function and MAX pooling layers (MAXPOOL) before a fully connected layer (FCON). For both of the convolution layers, kernel, stride, and padding sizes are determined as 5, 1, and 2, respectively, based on the performance. The distinctive features are selected at the max-pooling layers with a kernel and a stride size of 2, which reduce the dimensions from (32×32) to (8×8) at two passes.

In the second architecture, we used two consecutive 1D convolution layers (CONV1D) along with the ReLU activation function. A dropout (DROP) layer is added to alleviate the overfitting problem. Then, a max-pooling layer (MAXPOOL) and a fully connected layer (FCON) are added.

For both architectures, the softmax function is used as an activation function at the output layer and the Categorical Cross Entropy is used as a loss function. For parameters optimization, we used the Adam optimizer.

6 EVALUATION RESULTS

We performed experiments for each dataset, data representation, and threat model. We categorize the evaluations into three: Raw, Text-like, and Image-like.

Raw. First, we performed evaluations on the raw data. As the multi-class classification models require a fixed input shape and an arbitrary elevation profile does not have a specific length, we divided the elevation profiles into equal-length (32) chunks and use the raw data to train and test the models. For all datasets and threat models, we used a slightly modified version of the soft voting ensemble method while testing with raw data. In conventional ensemble learning techniques, the input is passed to different models and the final prediction is assigned based on the decisions of the models. However, in this study, we passed the equal-length chunks of a single input, i.e. elevation profile, to a single model, and then assigned the final prediction based on the decisions on the input chunks. Soft voting sums up the predicted probabilities for each class label for each input chunk and returns the class label with the highest probability as the final prediction.

Text-like. Second, we performed evaluations with text-like features: n-gram, and tf-idf. With n-gram features, we performed experiments using 10-fold cross-validation and by fixing the dimension of n-grams to 5 for all datasets and associated threat models. With tf-idf features, we performed 10-fold cross-validation and fixed the dimension of n-grams to 5 for all datasets and threat models.

The user-specific dataset contains overlapped and repetitive portions by nature. In the Simulations subsection, we simulated the same behavior on the mined datasets and performed the same evaluations for comparison.

Image-like. For the experiments on the image-like representations, we employed three methods in CNN: unweighted loss function, weighted loss function, and fine-tuning. In the unweighted and weighted loss function evaluations, we split the test data from the dataset by considering the sample size of the classes; we assigned probabilities for each class considering the inverse proportion to its size and then randomly selected test data with the associated probabilities. In fine-tuning evaluations, we performed 10-fold cross-validation in the last round where all the classes have the same sample size.

6.1 Raw and Text-like Data Evaluation

6.1.1 Direct Evaluation

❶ **Evaluating** TM-1**.** We trained and tested models with the user-specific dataset. As shown in Table 2, the user-specific dataset has an unbalanced sample size across classes. To mitigate bias, we use the same sample size for each class and change the number of classes at each step. The evaluation results are shown in Table 5. Due to the limited number of samples, the accuracy decreases as the number of classes increases. The only exception is C1D with n-grams and tf-idf. One-dimensional convolutions were able to capture the characteristics of elevation profiles even with a limited number of samples. The results show 99.80% accuracy with C1D, n-grams, and 4-class classification. With tf-idf and C1D, we obtained 99.00% and 99.94% accuracy with 3-class and 2-class classification, respectively.

LSTM gives better accuracy with raw data compared to the other text-like representations. LSTM performs better on the data where the ordering is decisive. With the text-like representations, the original ordering of the values is encoded separately, thus LSTM could not extract much information through the ordering.

For TM-1, RF also performs better with raw data. Since extracting features from the range and the ordering of the values are less demanding for decision trees, it is reasonable to observe such a pattern.

Other classification methods, i.e., SVM, MLP, and C1D, benefit more from the n-grams and tf-idf features.

Since the user-specific dataset is compiled from actual users, exhibiting mobility patterns, about 35% of the routes are overlapped. In a repetitive and overlapped setting, both training and testing splits may contain similar patterns leading to high accuracy scores. The results prove that a targeted attack on a person whose activity history is known will be successful with accuracy as high as 99.80%.

❷ **Evaluating** TM-2**.** While evaluating TM-2, the borough-level dataset is used. Individual models are created for

TABLE 5

The overall evaluation results for TM-1. Accuracy (%) with different data representations and classification models. In this table, the following abbreviations are used: SVM: Support Vector Machine; RF: Random Forest; MLP: Multi-Layer Perceptron; C1D: 1D Convolution; LSTM: Long Short Term Memory. **C** column indicates the number of classes in the classification problem. The following settings are used: 4-class = [WDC, ORL, NYC, SD], 3-class = [WDC, ORL, NYC], 2-class = [WDC, ORL].

C	raw data					n-grams					tf-idf				
	SVM	RF	MLP	C1D	LSTM	SVM	RF	MLP	C1D	LSTM	SVM	RF	MLP	C1D	LSTM
2	95.29	98.43	96.60	97.45	96.61	97.35	98.22	99.11	99.74	49.67	98.89	97.56	98.89	**99.94**	51.11
3	77.56	98.64	95.88	96.20	96.46	96.79	97.53	97.93	99.23	45.83	98.86	97.91	98.48	**99.00**	42.60
4	70.39	96.51	71.87	74.17	75.78	93.33	87.99	95.66	**99.80**	38.67	92.33	90.66	92.33	99.54	33.33

TABLE 6

The overall evaluation results for TM-2. Accuracy (%) with different data representations and classification models. In the table, we use the following abbreviations: LA: Los Angeles; MIA: Miami; NJ: New Jersey; NYC: New York City; SF: San Francisco; WDC: Washington, D.C.

Cities	raw data					n-grams					tf-idf				
	SVM	RF	MLP	C1D	LSTM	SVM	RF	MLP	C1D	LSTM	SVM	RF	MLP	C1D	LSTM
LA	67.18	77.41	63.43	32.50	36.88	**78.02**	74.33	77.27	55.29	28.25	76.27	76.02	75.33	58.49	23.00
MIA	69.57	80.85	67.17	80.55	54.10	75.55	77.55	75.77	77.33	40.33	83.77	82.44	72.00	**99.54**	25.33
NJ	65.56	82.56	78.31	83.32	74.18	74.76	66.19	82.39	71.19	39.05	77.93	82.69	65.71	**92.42**	32.14
NYC	73.63	**84.26**	73.44	37.33	25.57	82.25	80.71	79.78	73.02	20.72	81.50	82.96	82.63	76.60	17.94
SF	65.92	74.66	65.89	42.21	32.53	74.71	76.15	**80.25**	54.08	25.88	76.13	75.71	74.71	58.07	27.92
WDC	53.08	77.30	60.44	64.01	56.05	76.13	70.63	67.20	**89.61**	58.74	75.44	74.34	55.50	85.24	51.62

TABLE 7

The overall evaluation results for TM-3. Accuracy (%) with different data representations and classification models.

C	raw data					n-grams					tf-idf				
	SVM	RF	MLP	C1D	LSTM	SVM	RF	MLP	C1D	LSTM	SVM	RF	MLP	C1D	LSTM
3	61.53	**88.99**	64.29	82.53	56.15	76.70	78.04	77.20	65.23	33.92	85.22	81.33	81.81	65.99	33.20
5	73.14	**93.00**	73.53	65.65	70.11	80.33	78.67	79.53	53.92	20.28	86.18	82.78	84.11	52.38	20.11
7	77.58	**94.94**	80.60	52.58	64.98	85.22	84.73	84.82	43.50	13.86	88.59	88.01	87.27	48.14	14.23
8	80.60	**95.98**	80.50	44.57	67.55	84.77	84.85	85.12	43.79	14.03	87.19	86.24	86.55	50.28	12.50
10	83.72	**95.36**	84.11	40.81	59.20	87.46	87.78	87.12	31.22	16.95	89.48	87.99	88.41	43.86	12.56

each of the cities, by labeling the data as the name of the corresponding borough and evaluated separately. Similar to the user-specific dataset, the borough-level dataset also has an unbalanced sample size across the classes. To avoid biased results, we fix the sample size to that of the smallest class for all classes. At each fold, we randomly select train and test data for the classes with more samples. Table 6 shows the accuracy results of each model.

Los Angeles model reaches up to 78.02% accuracy with n-grams and SVM. Similar results are obtained with other classification methods and data representation pairs, such as raw and RF, n-grams and MLP, tf-idf, and SVM. For Los Angeles, C1D could not become prominent; the less complex models perform better on this dataset.

With **Miami**, we reach up to 99.54% accuracy with tf-idf and C1D. Overall, tf-idf is shown to be a better representation for this particular dataset. Combining a complex model with a representative feature, we achieved high accuracy.

For **New Jersey**, we achieve 92.43% accuracy with tf-idf and C1D. According to the results, tf-idf features better represent the New Jersey dataset.

In **New York City**, we reach up to 84.26% accuracy with raw data and RF. When we examine the dataset, we observed that the elevations fluctuate mostly between 13 ft and 95 ft. When such a small range is considered, decimal digit precision plays an important role. Since we do not

discard any precision in the raw dataset, it is reasonable to have better accuracy with raw data. Although the highest accuracy is obtained with raw data and RF, tf-idf is a better choice for other classification methods.

In **San Francisco**, we achieve 80.25% accuracy with n-grams and MLP. Both text-like representations present similar accuracy patterns.

For **Washington DC**, we obtain 89.61% accuracy with n-grams and C1D. For both text-like representations, C1D shows better performance than other methods.

Overall, we can clearly observe the difference between TM-1 results and TM-2 results. The two main reasons for this performance gap are that (i) there are no overlapped or repetitive routes among the mined segments in the borough-level dataset, and (ii) the elevation differences and elevation sequences are not distinctive enough within a city to decide in which borough is the given test data is. The results of the simulated behavior will be discussed in the simulations subsection.

❸ **Evaluating** TM-3. In TM-3 evaluations, due to sample size differences across the labels in the city-level dataset, we follow the same procedure in TM-1 evaluations. A fixed number of samples is randomly selected from each class for training and testing. Table 7 shows the results of the evaluation. Per the reported results, we are able to predict

the city of an elevation profile among 10 cities with an accuracy of 95.36%, among 8 cities with an accuracy of 95.98%, among 7 cities with an accuracy of 94.94%, among 5 cities with an accuracy of 93.00%, and among 3 cities with an accuracy of 88.99%. For TM-3, for all number of classes, we find the best performing configuration as raw data and RF. When we look into the reason for the fact that raw data with RF outperforms every other configuration, we observe that the elevation range of the different classes in this dataset plays an important role, similar to TM-2: NYC. Decision trees in RF are able to capture the features firsthand, without any representation needed in the middle.

When we consider the text-like representations, we observe that tf-idf features better represent the dataset. The success of the city-level estimations, when compared to the borough-level estimations (TM-2), is due to the elevation range and sequence differences across cities, which is reasonable, even though the dataset is mined in a similar fashion as in the borough-level dataset. This mining indicates that the city-level dataset also does not contain comprehensive, repetitive, and overlapped samples. The results of the simulated evaluation will be discussed in the following.

6.1.2 Simulations

The mined datasets do not contain overlapped or duplicate samples as in the user-specific dataset. In this set of evaluations, we simulate overlapped mined datasets and perform evaluations under the same threat models.

Simulation of TM-2. For the city-level estimation evaluations, we rebuild a simulation dataset with a $30 - 34\%$ overlap ratio for each region within the cities. The same evaluation procedures are then followed as the original mined dataset, which is 10-fold cross-validation with a fixed n-grams size of 5. 10(a) and 10(b) show the comparison between the best-achieved result in the original evaluation and the best-achieved result in the simulations. The increase in the accuracy confirms our previous hypothesis that having overlapped route samples would increase the accuracy. Since the mined dataset is not specific to any target user's mobility pattern, it is anticipated to result in less accuracy than the TM-1 evaluation accuracy scores.

Simulation of TM-3. For TM-3's simulated evaluations, we rebuild a simulation dataset with a 35% overlap ratio for each city and performed the same evaluation with 10-fold cross-validation and 5-grams. 10(c) shows the comparison of the best-achieved accuracy results in original evaluations and simulations. As expected, the accuracy is increased in the simulations proving our previous hypothesis that having similar patterns in a dataset affects the success of the attack.

6.2 Image-like Data Evaluations

In this set of evaluations, we perform experiments on the image-like representations of the data. Since the original data is unbalanced, the dataset built with the image-like representation also inherits the problem. In this section, we explain the methods to avoid bias due to an unbalanced dataset and discuss the associated results.

TABLE 8
Comparison of maximum achieved accuracy across different methods. The Unweighted Loss (UWL) column is not considered while deciding the maximum accuracy, as the results are biased. The maximum accuracy of each evaluation is written **bold**, the results that are not considered are written *italic*.

Methods	Raw	Text-like	Image-like		
		n-grams/ tf-idf	UWL *(biased)*	WL	FT
TM-1	96.51	**99.94**	*96.98*	95.23	87.93
TM-2: LA	77.41	**78.02**	*68.85*	68.39	63.63
TM-2: MIA	80.85	**99.54**	*88.96*	86.80	62.50
TM-2: NJ	83.32	**92.42**	*93.45*	79.42	57.14
TM-2: NYC	**84.26**	82.96	*74.20*	79.37	72.79
TM-2: SF	74.66	**80.25**	*67.20*	78.70	65.38
TM-2: WDC	77.30	**89.61**	*62.79*	70.28	71.50
TM-3	**95.36**	89.48	*92.51*	92.82	89.00

Dealing with Unbalanced Dataset. There are various methods to deal with unbalanced datasets, including downsampling, oversampling, and creating synthetic samples from existing ones. Among these methods, downsampling and oversampling are the easiest ones to explore, although downsampling leads to losing a great amount of data, and oversampling raises the chances of getting lower accuracy as the misclassified duplicated samples increase the false ratio. Therefore, we explore other alternatives: (i) weighted loss function and (ii) fine-tuning with different samples.

Weighted Loss Function. For the unbalanced dataset, we utilize a weighted loss function while training the CNN and use all the data in the dataset. By assigning a class weight that is inversely proportional to the sample size of the class, we signify samples of small classes while calculating the loss, thus their effect does not easily wear off.

Fine-Tuning with Different Samples. Fine-tuning is a common technique in deep learning and is used for re-training a complex pretrained model with another dataset. To address the unbalanced dataset, we take advantage of fine-tuning in a different manner. Namely, we introduce rounds and create a set of small datasets from the unbalanced datasets for each round. As illustrated in Figure 11, several small and balanced datasets are created by randomly selecting samples. For each consecutive round, samples of one or more classes are discarded, and the round dataset is created from the remaining classes. After round dataset creation, the model is trained with the round dataset that contains *the least number of classes*, *i.e.,* the lattermost created round dataset. At each step, the model is re-trained using the same or different hyperparameters until all the rounds expire. The dataset ordering of the rounds is reversed since the impact of the smallest dataset would wear off if the model is trained with the same order of round dataset creation, which conflicts with the whole idea. As illustrated in Figure 12, while re-training, the parameters of the previous model are passed to the model of the next round. The hyperparameters of each round can be tuned accordingly. For instance, for the last round, where we include all of the classes, the learning rate is reduced to find the loss minima.

To evaluate our attacks on the image-like data, the elevation profiles are converted into a dataset of images

(a) Simulation of TM-2 with n-gram features   (b) Simulation of TM-2 with tf-idf features   (c) Simulation of TM-3
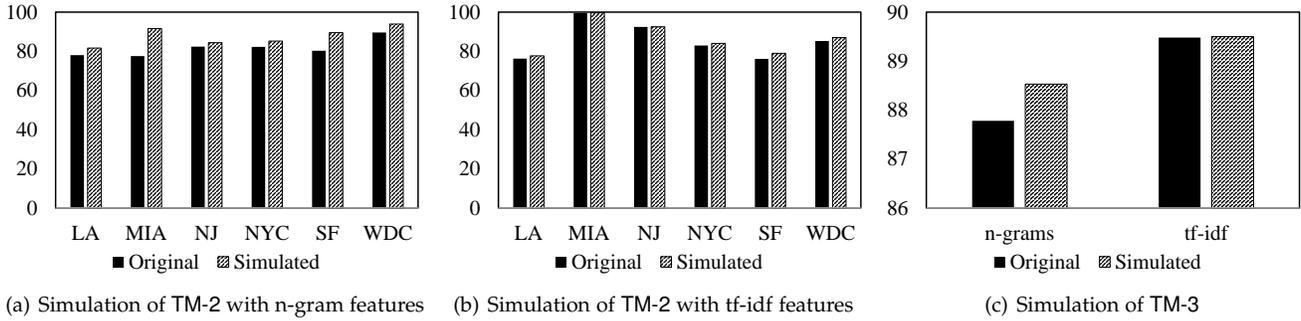
Fig. 10. Some selected simulation results of TM-2 and TM-3. The maximum achieved accuracy results are compared.
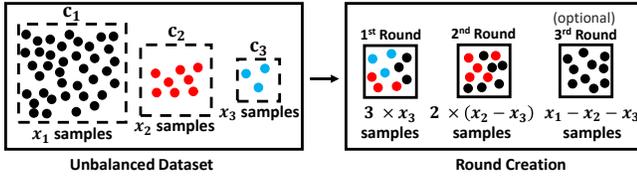
Fig. 11. An illustration of round creation from an unbalanced dataset of three classes.
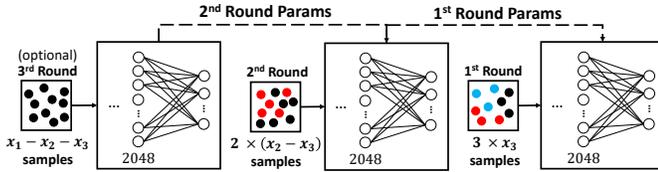
Fig. 12. An illustration of the fine-tuning pipeline for an unbalanced dataset of three classes.

TABLE 9
The fine-tuning results for TM-1 and TM-3 as the epoch size changes.

	TM-1			TM-3		
Epoch Size	**500**	**1000**	**2000**	**500**	**1000**	**2000**
Accuracy	79.31	87.96	82.73	86.04	89.00	87.85
Recall	55.87	67.54	63.12	29.76	45.34	38.91
Specificity	86.33	92.65	88.46	92.27	93.98	93.29
F1 Score	58.62	68.25	63.37	36.23	45.45	41.12

TABLE 10
The fine-tuning results for TM-2 as the epoch size is 1000 and learning rate is 0.001 for all rounds.

	LA	**MIA**	**NJ**	**NYC**	**SF**	**WDC**
Accuracy	63.61	62.52	57.13	72.84	65.34	71.55
Recall	28.02	25.66	40.03	18.15	30.76	73.27
Specificity	75.84	75.97	66.75	83.43	76.35	73.22
F1 Score	28.83	28.64	37.55	18.46	31.47	73.44

and rounds using the configurations and steps discussed above. Table 8 highlights the maximum achieved prediction accuracy along with comparisons with other methods.

Weighted vs. Unweighted Loss Function. To observe the impact of the weighted loss function, we conduct evaluations without giving any weight to the classes in the loss function while using an unbalanced dataset. We note that the unweighted loss function evaluation results are biased due to the unbalanced dataset. Table 8 shows the maximum achieved accuracy for each dataset and method. Even though the weighted loss function evaluation results are biased, which *seems* successful in outputing the largest class used during training and testing, the biased results remain behind 4 evaluations out of 8. In TM-1 and TM-3, the accuracy scores of unweighted and weighted loss functions are considerably close. Thus, we conclude that the weighted loss function improved the prediction performance primarily for TM-2.

Fine-tuning vs. Weighted Loss Function. For the fine-tuning evaluations, round datasets are created from the original data. For TM-1, with 4 classes, 3 rounds are created. For TM-3, with 10 classes, 5 rounds were created by eliminating 1, 2, 1, and 2 classes at each round, respectively. The dataset of TM-2 can be considered as a compilation of the dataset of 6 cities: Los Angeles (3 rounds), Miami (3

rounds), New Jersey (2 rounds), New York City (4 rounds), San Francisco (2 rounds), and Washington DC (1 round). Even though the main idea is to use all the data we have, we decided to downsample the classes with a large sample size. For instance, in the evaluation of TM-2: New York City, the biggest class has 5,455 samples whereas the second biggest class has 960 samples. In such cases, we did not create an additional round for only one class as this round would have a strong influence over the predictions, *i.e.,* overfitting.

Table 8 shows the fine-tuning method outperformed the weighted loss function method only for TM-2: WDC. The difference between the fine-tuning evaluation of Washington DC and others is that we were able to create only one round from the data. Overall, according to the results shown in Table 9 and Table 10, the fine-tuning evaluation is not as successful as the weighted loss function evaluation, since we still lose some data while creating rounds.

Text-like vs. Image-like Evaluations. When we compare text-like and image-like representations, we can conclude that text-like representation is a better choice for such attacks. For all evaluations except TM-3, text-like representation outperformed image-like representation. For TM-3 and TM-2:NYC, the raw data and RF configuration is the best choice.

7 DISCUSSION

Defenses. In this study, we are strictly concerned with the elevation profile as a representation of the location. We note that the location itself is the eventual modality of interest, but the exposed information to the adversary from which the adversary will make the inferences is the elevation. To that end, however, the same technique used for location perturbation [27] to defend against location privacy breaches can be applied on the elevation profile, although more straightforwardly and effectively. Namely, two broad classes of defenses could be applied to our problem domain: perturbation and aggregation.

Perturbation as a Defense. To thwart our inference attacks, we can perturb the elevation while maintaining its overall statistical features used for the original application by adding a carefully crafted Gaussian noise driven from the elevation distribution (zero mean and a standard deviation of the original data). We hypothesize that the noise will affect the classifier but will not impact the validity of some of the driven insight from the elevation for the user (e.g., total elevation[2]). To highlight the validity of this approach, we conduct a limited experiment to perturb 10% of the original raw data used in the original classification problem in Table 5. To keep the elevation signal plausible and convincing, we superimposed dependently generated noise over an epoch of time (10 seconds) and "clipped" the generated noise with one standard deviation of a moving average in the current segment[3]. As a result, we were able to reduce the accuracy from 95.29% (in the case of SVM, C=2) to 72.46%. Similar performance degradation is observed with the n-grams and tf-idf in the same settings: by generating those features from the perturbed data, we were able to achieve an accuracy of 67.19% and 70.31% with the n-grams and tf-idf, respectively. While not totally subverting the inference, the approach shows an initial promising direction.

Aggregation as a Defense. Another defense is realized by aggregating the elevation profile information into application-compliant statistics, e.g., total ascend, descend, mean ascend and descend, their standard deviation, minimum or maximum (over quantized elevation profile) or sampled elevation signal, which would reduce the effectiveness of our inference attack or block it altogether. The objective of the defense can be realized by rounding (10s of feet of elevation) so that the number of possible segments associated with the given statistical features is large enough to provide plausible deniability through anonymity (e.g., the correlation between the aggregate and location is weakened).

Feature-level Defenses and Caveats. We note that other approaches that directly perturb the feature representation modality (e.g., images [28]) may not be as practical, since the image itself is a constrained domain, and not every pixel

in the image domain is a valid perturbation candidate. We emphasize, however, that this issue is not particular to this problem space we address in this paper, but applicable to a range of problems in general, such as software [29]–[31] and network domains [32], where the feature representation used for implementing the machine learning algorithm transforms the input by upholding a dependency among the features, which is not the case in the original image modality used in computer vision applications [33].

Why Elevation and Actual Implications. We emphasize that we consider the elevation profile information as our inference input because this feature modality might be viewed, even to the most privacy-savvy individuals, as an innocent modality that does not necessarily expose much information about a specific location. Technically, however, it is not far-fetched to assume that there is an infinite number of mappings between a given elevation profile and location segments, which is shown to be not the case in this study under various plausible adversarial settings and objectives.

Moreover, we recognize the subjectivity of valuing location privacy by users [34], [35]. For instance, while some users might not care about sharing their exact location information all the time, some others may not feel comfortable sharing such information [36]. We demonstrate that such users, even when under the impression that they are not exposing their direct location, would be allowing an adversary to infer that location indirectly from the shared elevation. The sharing itself might put those users at risk, particularly when coupled with context. Imagine, for instance, an adversary who knows where a victim lives, but is able to infer that the activity of the victim, posted live, is associated with a location where the victim does not live, which would allow the adversary to stalk, or even break into the victim's house. By the same token, an adversary who is able to precisely locate such an activity, in real-time, might be able to launch the physical attack possible under the sharing of the direct location (e.g., theft of expensive biking gear [4]).

Other Activities and Associated Risks. We note that, in general, mobility patterns are both individual and activity-dependent [37]. For instance, the mobility of a salesperson would be totally different from the activity of a student. However, we note that our study is chiefly concerned with mobility patterns of specific activities: exercise. It is very difficult to envision a plausible scenario where a salesperson, for instance, would share the elevation profile of their activities moving door-to-door to sell a product. Similarly, while activities, in general, are tracked by various smartwatches (for everyday use), such tracking is limited to the high-level aggregate (e.g., total, sampled over epochs of time), and is not shareable directly. To this end, activities that are not exercise-related (on public trackable roads) are considered out of the scope of our use and attack models.

Explaining the Performance of Our Techniques. We note that, generally, the image and text representations outperformed the raw data used directly for our inference attack (except in the case of LSTM and random forest). The reason why LSTM performed well over the raw data is because of the natural mapping between such raw data (i.e., time series) and the operation of LSTM with attention to long

---

2. We emphasize that this defense will strictly preserve some but not all of the features of the elevation profile. For instance, by design, the perturbation will not preserve the total ascend and descend, two vital statistical features of the elevation profile.

3. We note that this plausibility step further restricts the perturbation to make the perturbed elevation profile acceptable by humans. This restriction explains the limited performance of the defense as unrestricted perturbation reduces the accuracy of the inference attack under the same settings (i.e., 10% perturbation, SVM, C=12) to 43.87%.

sequences and their dependencies. In other words, LSTM is capable of representing the features of time series well. On the other hand, the random forest technique is known to be tolerant to noise, so one plausible explanation for its superior performance is its implicit feature selection and representation (by creating decision points on ranges). On the other hand, CNN worked well with image-like representations for its power in the representation of features from such modality. By the same token, SVM and MLP are shown to be superior in the text-like features for their tolerance to noise achieved by adjusting decision margins.

In general, we also note that CONV1D performed better for TM-1, but not for TM-2 and TM-3, by demonstrating the power of CNN in general for capturing repeating patterns: TM-1 is concerned with profiling persons with a number of activities, some of which might be overlapping or even repeated, and CNN is known to extract high quality, representative, and discriminative features in that space, in contrast to the scenarios of TM-2 and TM-3, where such patterns are less manifested.

Finally, we note that natural languages in general have various semantic and syntactic characteristics that are not manifested in the non-constrained domain of elevation profiles. To that end, while LSTM, when applied to features driven from the tf-idf or n-grams of a natural language utterance, would perform well, the same technique is not guaranteed to perform as well using the same feature representations, as we demonstrate in this study. One plausible explanation also for the superior performance of LSTM over the raw data is that the raw data explicitly and fully maintains an ordering that is essential for the operation of LSTM, whereas that ordering is implicit and for very short utterances in tf-idf and n-grams. Losing such valuable information as a result of feature extraction would (although marginally) affect the performance, as shown in this study.

8 RELATED WORK

In this work, we addressed the problem of *location privacy* in activity trackers using the side-channel information obtained from publicly shared elevation profiles. While there is no work that explicitly addresses this topic, there is has some studies on various topics that are related in the broad literature [38]–[40]. In the following, we review some of those studies.

Most location privacy breaches are caused since users do not know why or how to preserve location privacy. [41] developed a tool to examine possible privacy exposures of users in their social networks where the data is mostly collected from wearable devices. Using this tool, the authors aimed to enhance the awareness of information leakage in social networks, particularly fitness apps in which the data retrieved from wearable devices is shared on social networks. Abdelmoty and Alrayes [42] aimed to increase awareness of location privacy on geo-social networks by surveying 186 users, where 77% of them indicated they use location-based services often, several times a day, and 47% of them reported that they were not aware that the location-based apps collect and store location information even when users select the private location option. Moreover, 43% of

respondents were not aware that applications may share location information with third parties.

Despite the methods employed to preserve location privacy, several attacks are devised to uncover supposedly protected locations. Experiments for revealing exact locations from trajectories with private zones are conducted on a fitness-tracking social network, Strava [18]. Researchers found the exact endpoints associated with users, even when such users selected the private zone option when sharing the training route. In another study, location trajectories of users are recovered from publicly available aggregated mobility data obtained from GSM operators [43]. The attack relies on tracking the regularity—i.e., coming across the same location trace in the aggregated data regularly—and uniqueness—i.e., the location trace belongs to a unique user—of the user mobility traces to recover trajectories.

As our study exemplifies, online social networks lay under the scope of privacy breach risks for users. Zheng *et al.* [44] shows that sharing data that reveals spatiotemporal features of users' mobility patterns on online social networks reveals sensitive information such as home location, using a different form of data, *i.e.,* multimedia. Rossi *et al.* [45] show that location-based social networks are vulnerable to identity privacy breaches by revealing the identity of users by observing their mobility patterns.

Several attacks against general location privacy methods are proposed [46]. The homogeneity attack [47] is an attack on k-anonymity to infer data of interest from other shared data. Machanavajjhala *et al.* [47] illustrated a scenario where an adversary infers the illness of a target person from available information, the zip code, age, etc. The same method can be applied to infer location data. In location distribution attacks [48], the adversary exploits the fact that users are mostly not uniformly distributed in the location space. Another attack by Shokri *et al.* [49] utilized the aggregated traffic statistics and environmental context information. The attack scenario includes an adversary who tries to reveal the possible location of the target by making use of the fact that the probability of the target's whereabouts is not uniformly distributed. Map matching methods [50] aim to restrict the obfuscated area to a smaller but plausible area by removing irrelevant areas. Movement boundary attacks were explored [51], where the adversary aims to calculate the movement boundary of a target by chasing the position queries and updates of the target. After calculating the boundary, the location of interest, such as home or workplace, is inferred and the irrelevant locations are discarded.

Although we did not directly touch upon preserving the location privacy in our study, there have been a few related studies in this space. The fast-growing need of preserving location privacy over the aforementioned attacks excited researchers' attention. Researchers introduce obfuscation methods such as decreasing the quality of the location by introducing inaccuracy and imprecision [52]. Additionally, the term k-anonymity is defined as obscuring the location information of individuals with *k* number of other individuals within the region [53], [54].

9 CONCLUSION

In this paper, we presented new attacks on location privacy using only elevation profiles. The attacks are categorized

into three types: predicting location by knowing the activity history of the target, predicting the borough by knowing the city of the target, and predicting the city of the target without any prior knowledge. The key contributions of our work are proving the concept that hiding the route of a workout and sharing only the elevation profile is not sufficient to preserve location privacy, defining a new attack surface by creating scenarios for possible threat models, and providing a machine-learning approach to realize such threat as attacks. To validate our attacks we created three datasets by collecting data from athletes and mining data from a popular fitness-tracking website and Google Elevation API. We preprocessed the datasets by employing Natural Language Processing and Computer Vision approaches and then employed classification techniques to predict the location from elevation profiles. En route, we defined three threat models and evaluated each of them individually on the different datasets. As a result of the evaluations, we were able to identify the corresponding location of an elevation profile with accuracy between 59.59% and 95.83%.

While this work highlights the clear trade-offs provided by the various defenses, their usability is largely unexplored. In our future work, we will explore the usability of compatible defenses such as devising and using route statistics that serves the same purpose as sharing elevation profile – demonstrating the roughness of the route, while preserving users' privacy and acceptance.

REFERENCES

[1] Ü. Meteriz, N. F. Yiotaldiotaran, J. Kim, and D. A. Mohaisen, "Understanding the potential risks of sharing elevation information on fitness applications," in *40th IEEE International Conference on Distributed Computing Systems, ICDCS*, 2020, pp. 464–473. [Online]. Available: https://doi.org/10.1109/ICDCS47774.2020.00063

[2] J. P Higgins, "Smartphone applications for patients? health & fitness," *The American journal of medicine*, vol. 129, 06 2015.

[3] I. Polakis, G. Argyros, T. Petsios, S. Sivakorn, and A. D. Keromytis, "Where's wally?: Precise user discovery attacks in location proximity services," in *Proceedings of the 22Nd ACM SIGSAC Conference on Computer and Communications Security*, ser. CCS '15. New York, NY, USA: ACM, 2015, pp. 817–828. [Online]. Available: http://doi.acm.org/10.1145/2810103.2813605

[4] G. Friedland and R. Sommer, "Cybercasing the joint: On the privacy implications of geo-tagging," 08 2010, pp. 1–8.

[5] "Map privacy," https://bit.ly/2TgWTJu, accessed: 2020-01-11.

[6] "How do i remove display map from my activity list?" https://bit.ly/2TPYdDf, accessed: 2020-01-11.

[7] "Hide map," https://bit.ly/3gaqKwe, accessed: 2020-01-11.

[8] "Privacy setting to hide activity map from non-followers," https://bit.ly/3cu4ET9, accessed: 2020-01-11.

[9] S. Padilla, I. Mujika, G. Cuesta, and J. Goiriena, "Level ground and uphill cycling ability in professional road cycling." *Medicine and science in sports and exercise*, vol. 31 6, pp. 878–85, 1999.

[10] "Strava global heatmap," https://www.strava.com/heatmap, accessed: 2019-03-20.

[11] "U.s. soldiers are revealing sensitive and dangerous information by jogging," https://goo.gl/tiM5VU, accessed: 2021-04-20.

[12] "Fitness tracking app strava gives away location of secret us army bases," https://bit.ly/3vdfVNT, accessed: 2021-04-20.

[13] S. Loughran, "Advanced deanonymization through strava," https://bit.ly/3cNURHN, 2019, accessed: 2021-04-20.

[14] "Strava hone app led thieves to my 12,000 bike collection," https://bit.ly/354wcKB, accessed: 2021-04-20.

[15] "Cyclists warned to beware sharing data on ride-tracking apps," https://bit.ly/3g66fRe, accessed: 2021-04-20.

[16] "Cyclist who had five bikes stolen says thieves are looking for quick times on strava to try and find high-end bikes ? warns other users to check their privacy settings," https://bit.ly/3gp8xtp, accessed: 2021-04-20.

[17] "Organised criminals using fitness apps to track expensive bicycles and equipment," https://bit.ly/2SlR11n, accessed: 2021-04-20.

[18] W. U. Hassan, S. Hussain, and A. Bates, "Analysis of privacy protections in fitness tracking social networks -or- you can run, but can you hide?" in *27th USENIX Security Symposium (USENIX Security 18)*. Baltimore, MD: USENIX Association, 2018, pp. 497–512. [Online]. Available: https://www.usenix.org/conference/usenixsecurity18/presentation/hassan

[19] "Runtastic deserves more ?heat? than strava!" https://bit.ly/2SqtCvD, accessed: 2021-04-20.

[20] I. Maglogiannis, K. Karpouzis, M. Wallace, and J. Soldatos, Eds., *Emerging Artificial Intelligence Applications in Computer Engineering - Real Word AI Systems with Applications in eHealth, HCI, Information Retrieval and Pervasive Technologies*, ser. Frontiers in Artificial Intelligence and Applications, vol. 160. IOS Press, 2007.

[21] C.-W. Hsu, C.-C. Chang, and C. Lin, "A practical guide to support vector classication," 2008.

[22] L. Breiman, "Arcing classifiers," *ANNALS OF STATISTICS*, vol. 26, pp. 801–823, 1998.

[23] ——, "Random forests," *Mach. Learn.*, vol. 45, no. 1, pp. 5–32, Oct. 2001. [Online]. Available: https://doi.org/10.1023/A:1010933404324

[24] T.-S. Lim, W.-Y. Loh, and Y.-S. Shih, "A comparison of prediction accuracy, complexity, and training time of thirty-three old and new classification algorithms," *Machine Learning*, vol. 40, no. 3, pp. 203–228, Sep 2000. [Online]. Available: https://doi.org/10.1023/A:1007608224229

[25] P. W. Eklund, "A performance survey of public domain supervised machine learning algorithms," Tech. Rep., 2002.

[26] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," *CoRR*, vol. abs/1412.6980, 2014. [Online]. Available: http://arxiv.org/abs/1412.6980

[27] R. Dewri, "Local differential perturbations: Location privacy under approximate knowledge attackers," *IEEE Transactions on Mobile Computing*, vol. 12, no. 12, pp. 2360–2372, 2012.

[28] S. J. Oh, M. Fritz, and B. Schiele, "Adversarial image perturbation for privacy protection a game theory perspective," in *2017 IEEE International Conference on Computer Vision (ICCV)*. IEEE, 2017, pp. 1491–1500.

[29] H. Alasmary, A. Khormali, A. Anwar, J. Park, J. Choi, A. Abusnaina, A. Awad, D. Nyang, and A. Mohaisen, "Analyzing and detecting emerging internet of things malware: A graph-based approach," *IEEE Internet of Things Journal*, vol. 6, no. 5, pp. 8977–8988, 2019.

[30] A. Khormali, A. Abusnaina, S. Chen, D. Nyang, and D. Mohaisen, "From blue-sky to practical adversarial learning," in *Second IEEE International Conference on Trust, Privacy and Security in Intelligent Systems and Applications, TPS-ISA*, 2020, pp. 118–127.

[31] H. Alasmary, A. Abusnaina, R. Jang, M. Abuhamad, A. Anwar, D. Nyang, and D. Mohaisen, "Soteria: Detecting adversarial examples in control flow graph-based malware classifiers," in *40th IEEE International Conference on Distributed Computing Systems, ICDCS*, 2020, pp. 888–898.

[32] A. Chernikova and A. Oprea, "Fence: Feasible evasion attacks on neural networks in constrained environments," *ACM TOPS*, 2022.

[33] A. Abusnaina, Y. Wu, S. Arora, Y. Wang, F. Wang, H. Yang, and D. Mohaisen, "Adversarial example detection using latent neighborhood graph," in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2021, pp. 7687–7696.

[34] K. Fawaz, H. Feng, and K. G. Shin, "Anatomization and protection of mobile {Apps?} location privacy threats," in *24th USENIX Security Symposium (USENIX Security 15)*, 2015, pp. 753–768.

[35] B. P. Knijnenburg, A. Kobsa, and H. Jin, "Preference-based location sharing: are more privacy options really better?" in *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, 2013, pp. 2667–2676.

[36] J. Krumm, "A survey of computational location privacy," *Personal and Ubiquitous Computing*, vol. 13, no. 6, pp. 391–399, 2009.

[37] S. Hasan, X. Zhan, and S. V. Ukkusuri, "Understanding urban human activity and mobility patterns using large-scale location-based data from online social media," in *Proceedings of the 2nd*

ACM SIGKDD international workshop on urban computing, 2013, pp. 1–8.

[38] S. Narain, A. Ranganathan, and G. Noubir, "Security of gps/ins based on-road location tracking systems," in *2019 IEEE Symposium on Security and Privacy (SP)*. IEEE, 2019, pp. 587–601.

[39] S. Narain, T. D. Vo-Huu, K. Block, and G. Noubir, "The perils of user tracking using zero-permission mobile apps," *IEEE Security & Privacy*, vol. 15, no. 2, pp. 32–41, 2017.

[40] S. Narain and G. Noubir, "Mitigating location privacy attacks on mobile devices using dynamic app sandboxing," *Proceedings on Privacy Enhancing Technologies*, vol. 2019, no. 2, pp. 66–87, 2019.

[41] A. Aktypi, J. Nurse, and M. Goldsmith, "Unwinding ariadne's identity thread: Privacy risks with fitness trackers and online social networks." Association for Computing Machinery, 2017, pp. 1–11.

[42] A. I. Abdelmoty and F. Alrayes, "Towards understanding location privacy awareness on geo-social networks," *ISPRS International Journal of Geo-Information*, vol. 6, no. 4, 2017. [Online]. Available: http://www.mdpi.com/2220-9964/6/4/109

[43] Z. Tu, F. Xu, Y. Li, P. Zhang, and D. Jin, "A new privacy breach: User trajectory recovery from aggregated mobility data," *IEEE/ACM Transactions on Networking*, vol. 26, no. 3, pp. 1446–1459, June 2018.

[44] D. Zheng, T. Hu, Q. You, H. A. Kautz, and J. Luo, "Towards lifestyle understanding: Predicting home and vacation locations from user's online photo collections," in *Proceedings of the Ninth International Conference on Web and Social Media, ICWSM 2015, University of Oxford, Oxford, UK, May 26-29, 2015*, 2015, pp. 553–561. [Online]. Available: http://www.aaai.org/ocs/index.php/ICWSM/ICWSM15/paper/view/10487

[45] L. Rossi, M. J. Williams, C. Stich, and M. Musolesi, "Privacy and the city: User identification and location semantics in location-based social networks," in *Proceedings of the Ninth International Conference on Web and Social Media, ICWSM 2015, University of Oxford, Oxford, UK, May 26-29, 2015*, 2015, pp. 387–396. [Online]. Available: http://www.aaai.org/ocs/index.php/ICWSM/ICWSM15/paper/view/10498

[46] M. Wernke, P. Skvortsov, F. Dürr, and K. Rothermel, "A classification of location privacy attacks and approaches," *Personal and Ubiquitous Computing*, vol. 18, no. 1, pp. 163–175, Jan 2014. [Online]. Available: https://doi.org/10.1007/s00779-012-0633-z

[47] A. Machanavajjhala, J. Gehrke, D. Kifer, and M. Venkitasubramaniam, "L-diversity: privacy beyond k-anonymity," in *22nd International Conference on Data Engineering (ICDE'06)*, April 2006, pp. 24–24.

[48] M. F. Mokbel, "Privacy in location-based services: State-of-the-art and research directions," in *2007 International Conference on Mobile Data Management*, May 2007, pp. 228–228.

[49] R. Shokri, G. Theodorakopoulos, J. Le Boudec, and J. Hubaux, "Quantifying location privacy," in *2011 IEEE Symposium on Security and Privacy*, May 2011, pp. 247–262.

[50] J. Krumm, "Inference attacks on location tracks," in *Pervasive Computing*, A. LaMarca, M. Langheinrich, and K. N. Truong, Eds. Berlin, Heidelberg: Springer Berlin Heidelberg, 2007, pp. 127–143.

[51] G. Ghinita, M. L. Damiani, C. Silvestri, and E. Bertino, "Preventing velocity-based linkage attacks in location-aware applications," in *Proceedings of the 17th ACM SIGSPATIAL International Conference on Advances in Geographic Information Systems*, ser. GIS '09. New York, NY, USA: ACM, 2009, pp. 246–255. [Online]. Available: http://doi.acm.org/10.1145/1653771.1653807

[52] M. Duckham and L. Kulik, "A formal model of obfuscation and negotiation for location privacy," in *Proceedings of the Third International Conference on Pervasive Computing*, ser. PERVASIVE'05. Berlin, Heidelberg: Springer-Verlag, 2005, pp. 152–170. [Online]. Available: http://dx.doi.org/10.1007/11428572_10

[53] L. Sweeney, "Achieving k-anonymity privacy protection using generalization and suppression," *Int. J. Uncertain. Fuzziness Knowl.-Based Syst.*, vol. 10, no. 5, pp. 571–588, Oct. 2002. [Online]. Available: http://dx.doi.org/10.1142/S021848850200165X

[54] A. Gkoulalas-Divanis, P. Kalnis, and V. S. Verykios, "Providing k-anonymity in location based services," *SIGKDD Explor. Newsl.*, vol. 12, no. 1, pp. 3–10, Nov. 2010. [Online]. Available: http://doi.acm.org/10.1145/1882471.1882473

Ülkü Meteriz-Yıldıran (S'20) obtained her Ph.D. and M.Sc. degrees from the University of Central Florida, in Computer Science with a Machine Learning concentration, in 2022 and 2022. She earned her B.Sc. degree in Computer Engineering from Middle East Technical University in 2018. Her research interests are in applied machine learning on the privacy and security of wearable devices. She is a member of the Security and Analytics Lab (SEAL) at the University of Central Florida since 2018. In 2021, she interned at Amazon Web Services (AWS) as a software development engineer and is currently a machine learning researcher at Meta.

Necip Fazıl Yıldıran obtained his Ph.D. from the University of Central Florida in 2022 with a research focus on the analysis of highly configurable software. He earned his M.Sc. in Computer Science with a concentration on applied machine learning on the security and privacy of IoT, also from the University of Central Florida, in 2020. He earned his B.Sc. degree in Computer Engineering from Middle East Technical University in 2018. He interned at Google as a Software Engineer in 2020 and 2021, where he is currently a full-time Software Engineer.

Joongheon Kim (M'06–SM'18) has been with the School of Electrical Engineering, Korea University, Seoul, Korea, since 2019, where he is currently an associate professor. He received the B.S. and M.S. degrees in computer science and engineering from Korea University, Seoul, Korea, in 2004 and 2006, respectively; and the Ph.D. degree in computer science from the University of Southern California (USC), Los Angeles, CA, USA, in 2014. Before joining Korea University, he was with LG Electronics (Seoul, Korea, 2006–2009), InterDigital (San Diego, CA, USA, 2012), Intel Corporation (Santa Clara in Silicon Valley, CA, USA, 2013–2016), and Chung-Ang University (Seoul, Korea, 2016–2019). He serves as an associate editor for *IEEE Transactions on Vehicular Technology*. He published more than 90 journals, 110 conference papers, and 6 book chapters. He also holds more than 50 granted patents.

David Mohaisen (M'06-SM'15) earned his M.Sc. and Ph.D. degrees from the University of Minnesota in 2011 and 2012, respectively. He is a Full Professor of Computer Science at the University of Central Florida, where he leads the Security and Analytics Lab (SEAL) and has been since 2017. Previously, he was an Assistant Professor at SUNY Buffalo (2015-2017) and a Senior Scientist at Verisign Labs (2012-2015). His research interests are in the broad area of applied security and privacy, covering aspects of computer and networked systems, software systems, IoT and AR/VR, and machine learning. His research has been supported by NSF, NRF, AFRL, AFOSR, etc., and has been published in top conferences and journals alike, with multiple best paper awards. His work was featured in the New Scientist, MIT Technology Review, ACM Tech News, Science Daily, etc. Among other services, he is currently an Associate Editor of IEEE Transactions on Mobile Computing and IEEE Transactions on Parallel and Distributed Systems. He is a senior member of ACM (2018) and IEEE (2015).