

Guess Which Car Type I Am Driving: Information Leak via Driving Apps

Dongyao Chen
Shanghai Jiao Tong University
chendy@sjtu.edu.cn

Mert D. Pesé
Clemson University
mpese@clemson.edu

Kang G. Shin
University of Michigan, Ann Arbor
kgshin@umich.edu

Abstract—Driving apps, such as navigation, fuel-price, and road services, have been deployed and used widely. The car-related nature of these services may motivate them to infer the type of their users' vehicles. We first apply systematic analytics on real-world apps to show that the vehicle-type — seemingly unharmed — information may have serious privacy implications. Next, we demonstrate that attackers can harvest the features of these mobile apps to infer the car-type information in a stealthy way. Specifically, we explore the use of zero-permission mobile motion sensors to extract spectral features for differentiating the engines and body types of vehicles. Based on our experimental results of 17 different cars, we have achieved 82+% and 85+% overall accuracy in identifying three major engine types and four popular body types, respectively.

I. INTRODUCTION

This paper presents a new finding: the *motion sensors used by driving apps* on commodity smartphones may reveal the car-type information, thus posing a severe threat to users' privacy. This finding is motivated by two essential pieces of background information: 1) the profound implication of car-type information, and 2) the excessive and unregulated usage of zero-permission motion sensors in driving apps.

Car-type information represents the physical configuration — e.g., engine and body type — of the car. This data has intrinsic connections to highly private information, including spending habit, demographics, and even political leaning, e.g., conservative vs. liberal. For example, social perceptions [1] and prior research [2] have shown that the vehicle body type has a strong correlation with the owner's personal information [3]. This rich implication of a vehicle one owns/drives can, therefore, be an attractive target for various parties — including advertisers, political campaigners, lobbyists, etc. — who may exploit and monetize it. Adversaries may harvest the vehicle characteristics and then exploit it to infer people's sensitive information, thus sabotaging their privacy. For example, based on the correlation between people's ideology and cars they drive [2], the adversary can infer victims' most probable political tendencies, thus enabling various adversarial practices [4].

Due to the sensitivity of the characteristics of people's vehicle, its public access is usually restricted. For example, in the U.S., people's vehicle information data is organized & maintained by Department of Motor Vehicles (DMV) — a designated government agency, and is accessible *only* by authorized entities, e.g., the law enforcement. A natural question

is then “is it possible to mitigate the hindrance of accessing the user's car-type information?”

We explore if the car-type information can be leaked via motion sensors being used by driving apps, such as navigation (Google maps, Waze), driver services (GasBuddy and ParkWhiz), and dashcam, to name a few. We focus on driving apps because they usually have close connection to their user's driving, thus offering opportunities for attackers to implement practical side-channel privacy attacks on the car-type information. The usage of motion sensors has become essential in driving apps for a wide range of tasks. For example, navigation apps use the accelerometer and dead-reckoning to track moving vehicles when the satellite signal is unavailable (e.g., in a tunnel). For driver assistance apps, motion sensors are widely used for detecting abrupt braking and/or acceleration. However, despite their ever-increasing popularity, motion sensors do not require any privacy check, e.g., the user's consent. These features together make the motion sensor an ideal data source for achieving the side-channel information leak. An important question is then “is it possible to infer the vehicle-type information with the motion sensor?”

Our proposed system, called *VeFi*, is the first that enables car-type information to be leaked via commodity driving apps with zero-permission. In particular, we will show that by using the motion sensors on commodity smartphones, one can classify a vehicle's engine type (e.g., number of cylinders) and body type (e.g., compact, sedan, SUV, and pick-up truck). Hence, attackers can leverage *VeFi* to drastically *lower* the bar for accessing/infering the user's vehicle characteristics and thus causing large-scale leakage of various private information, such as people's spending habits, demographic information, and even political affiliations.

Characterizing vehicles based on their vibration patterns is, however, challenging because of the complexity of the vehicle vibration and the simplicity of motion sensors. Specifically, a vehicle is a complex compound of mechanical and electrical components. A vehicle's vibration pattern varies greatly with its physical state, e.g., idling (the vehicle is stationary while the engine is on) vs. moving. Such a variability makes it difficult to develop/use a unified approach to characterizing the vehicle's vibration pattern. Moreover, motion sensors generate a time-series data that only provides limited information. For example, the common sampling rate of motion sensors is 100 Hz, making it impossible to capture a signal with the vibration frequency above 50Hz, according to the Nyquist Theorem [5].

To overcome these challenges with the motion sensors of commodity mobile devices, we propose a novel scheme with flexible modality for characterizing vehicles in different states. Specifically, for an idling vehicle, we found the spectral feature of a high-frequency component can lead to insights of the vehicle's

engine-type. Note that our idling vehicle’s analytics pipeline is *training-less* — one (e.g., a malicious entity) can characterize the engine-type based on the sensor readings and then derive insights without training a statistical model. To characterize moving vehicles, we propose a new feature engineering scheme that can capture the inherent vibration patterns of different vehicle body-types. The thus-extracted feature can be applied to various machine-learning classifiers for determining the vehicle’s *body-type*. Using the real-world driving data collected from 17 different vehicles, our training-less scheme achieves 82+% overall accuracy in classifying engines with different numbers of cylinders. The data analytics pipeline for moving vehicles achieves 85+% accuracy in differentiating four representative vehicle types.

II. BACKGROUND

A. Primer of Motion Sensors

The motion sensors embedded in a mobile device are to capture its user’s motion. Their versatile roles in motion sensing (i.e., angular speed and acceleration) is achieved by the embedded inertial measurement units (IMUs). According to a recent survey [6], the penetration rates of accelerometer and gyroscope in modern smartphones are 100% and 48%, respectively. Due to their seemingly low privacy sensitivity, mobile OSs do not require the user’s explicit permission for apps’ access, *c.f.*, location information. In fact, the motion sensors are usually *always-on*, meaning that even if the user switches it to operate in the background, the data collection continues. These features together allow adversaries to collect the motion sensors data freely and stealthily without the victim’s notice.

B. Motion Sensor as a Vibration Sensor

Modern IMU sensors, often integrated into a compact MEMS chip in a mobile/wearable device, are capable of quantifying the device’s angular speed with time-series data. For example, the MEMS gyroscope measures the angular speed by harvesting the Coriolis force [7] when the device is rotating. To capture the Coriolis force, the gyroscope uses a tuning fork configuration [8] — a built-in fork-shaped mass oscillates at a high frequency. The design of the MEMS chip has led us to a key observation: *the oscillating mass can be used as a sampler for capturing the vibration of other objects*. Specifically, vibrations can affect the oscillatory pattern of the moving mass of the MEMS chip, which manifests itself as motion sensor readings.

C. Implications of Vehicle Characteristics

Vehicles are categorized by their physical configuration, such as engine- and body-type. In this paper, we will demonstrate the feasibility of differentiating three major engine types (i.e., 4, 6, and 8-cylinder engines) and four popular and representative body-types of vehicles (i.e., compact, mid-size, SUV, and pickup trucks).

The vehicle’s physical characteristics (especially the body-type) can imply a wide range of personal information, according to social science research [2], market study [9], and public perceptions [1]. Hence, the vehicle-characteristics information is prohibited from public access, e.g., in the US, this information is maintained by each state’s Department of Motor Vehicles (DMV) or Secretary of State (SOS). For example, since 4-cylinder engines are less powerful and have higher miles-per-gallon

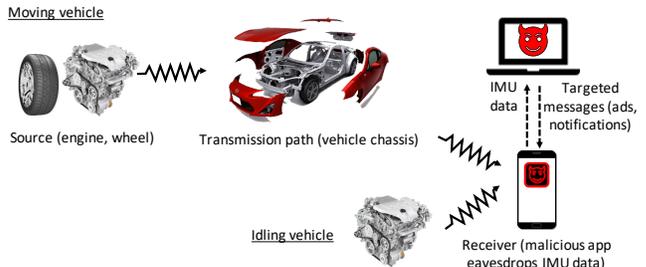


Fig. 1. Overview of attack models in VeFi.

(MPG) than 8-cylinder engines, the engine type information may reflect the vehicle owner’s life style and preference in fuel consumption. In what follows, we use real-world examples to show that several unique privacy-sensitive implications can be inferred from the vehicle-type information.

Spending habit. People’s choice of vehicle has close association with their spending habit. This observation has already been monetized by targeted advertising companies [10].

Political preferences. Prior research [11], [12], market study [9], and surveys [1], [13], [14] indicate that people’s choice of vehicle may also reflect their political leaning (e.g., conservative vs. liberal). Such information is highly sensitive, even constituting an attractive target for adversaries to influence people’s voting in an election.

Demographic information. Some specific demographic information (e.g., gender and education level) also turns out to be closely associated with people’s choice of vehicle. For instance, the Gallup survey [14] studied the most frequently driven cars by Americans. The results indicate a strong correlation between people’s education level with choice of their cars in the U.S. — only 39% of the participants with postgraduate education drive large US-made cars, while about 61% of those drive cars that are made by foreign companies (e.g., Toyota and BMW).

Even though vehicle characteristics may not be the most explicit indicator for the aforementioned personal attributes, with such information from a large group of people, an adversarial entity may accumulate the side-channel information to effectively pose a large-scale privacy threat. A notable example was the Cambridge Analytica scandal [4]. Specifically, the data analytics company, Cambridge Analytica, managed to obtain approximately 87 millions Facebook users’ social network activity data (e.g., Facebook-likes and shares). Based on this non-explicit information, Cambridge Analytica managed to infer people’s political tendency, which was then used to influence the people’ ideology by fabricating fake news and distributing them to targeted victims. This incident was reported to have influenced the 2016 US presidential election [15].

III. THREAT MODEL

In VeFi, the victims are the users of a mobility app that stealthily collects IMU sensor data for vehicle characterization. We also assume each victim owns and drives a car while using the app that collects IMU data. This is a reasonable assumption since the vehicle is one of the most popular personal properties (e.g., in the U.S., 95% of households own a car, 85% of them get to work by cars [16]), whereas most smartphones are

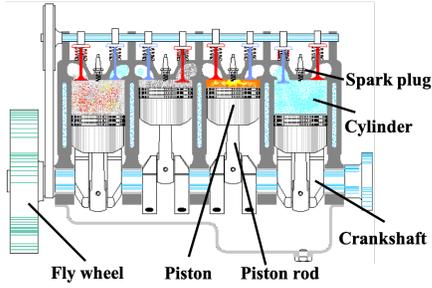


Fig. 2. Illustration of IC engine powertrain [18].

equipped with IMU sensors. We will later justify the use of smartphones for the detection of a driving event.

As shown in Fig. 1, the attackers can be adversarial entities that have access to the victim’s driving data, including malicious app developers who publish smartphone apps that eavesdrop the IMU data for characterizing the users’ vehicle and/or third-party data processor who are interested in monetizing the information of end-users’ vehicles. VeFi enables a practical and large-scale side-channel privacy leakage attack even for those attackers with limited resources (i.e., zero-permission IMU data). For example, if the attacker’s goal is to send targeted ads, s/he only needs to infer the victim’s vehicle configuration and present the corresponding ads in the app’s advertising banner.

IV. CHARACTERIZING IDLING VEHICLES

To characterize a vehicle’s vibration pattern, it is essential to understand internal combustion (IC) engine dynamics and the vibration induced by the IC engine.

A. Fundamentals of IC Engine

We focus on traditional vehicles with gasoline engines, which is a subcategory of IC engines. The US Energy Information Administration reports the share of gasoline-powered vehicles in 2018 among new sales to be 93%, and predicts the share of gasoline and flex-fuel vehicles (which use gasoline blended with up to 85% ethanol) to be 75% in 2050 [17]. As a result, gasoline engines will remain as the dominant type of propulsion for the decades to come.

Overview of the IC Engine. To convert the chemical energy (i.e., from gasoline) to mechanical power, each cylinder of the IC engine (as shown in Fig. 2) combusts the vaporized fuel to generate force to move the piston linearly. The linear movement is then converted to rotational motion of the crankshaft, which rotates the fly wheel and powers up the vehicle.

Engine Working Cycle. The engine vibration is induced by the reciprocating working cycle [19] that is adopted in the vast majority of IC engines. Specifically, each engine completes four strokes (i.e., intake → compression → combustion → exhaust) to turn the crankshaft a full rotation of 360° . Note that in one crankshaft revolution, only a half of the engine cylinders are sparked sequentially. That is, in each revolution, there are $\frac{C}{2}$ combustions in the engine, where C is the number of cylinders of the engine. For example, for a 4-cylinder engine, one revolution generates two fuel combustions, which are the main cause of engine vibration.

Vibration Frequency of IC Engine. The engine’s vibration frequency can be inferred from its structure and speed, measured

in revolutions per minute (RPM), or the number of crankshaft revolutions per minute. Specifically, for the reciprocating IC engine, we can get the combustion frequency:

$$f_C = \frac{RPM}{60} \cdot \frac{C}{2} \quad (1)$$

B. Fundamental Frequency Analysis

For a vibrating object, the detected vibrating frequency is a composite of several frequency components, i.e., fundamental frequency and harmonics [20]. The fundamental frequency is the component that has the lowest frequency. Harmonics are multiples of the fundamental frequency, with lower magnitudes than the fundamental frequency. For an idling vehicle, the harmonics induced by the engine can be represented as:

$$f_{C,N} = N f_C \quad (2)$$

where N is a positive integer, $N \in \mathbb{N}$; $f_{C,n}$, $n \geq 1$ is the n^{th} -order harmonic; $f_{C,1}$ is the fundamental frequency. The vibration that induces $f_{C,n}$ is called n^{th} -order vibration. Note that the magnitude induced by the 1st-order vibration dominates the whole-body vibration of an idling vehicle. According to the survey [21] on the vast majority of IC engine-types (i.e., 4, 6, and 8-cylinder), the 1st-order vibration contributes to more than 70% of the whole-body vibration of an idling vehicle.

C. Frequency Aliasing

One of the key limitation of gyroscope sensor is its narrow bandwidth — gyroscope may not have a sufficient sampling rate to capture the vehicle’s vibration. Specifically, according to the Shannon-Nyquist Theorem [5], the sampling rate needs to be at least twice the target’s frequency. That is, given the sampling rate of the gyroscope as f_s , the gyroscope cannot capture $f_{C,N}$ if $f_{C,N} > \frac{1}{2}f_s$. To overcome this problem, we harvest the *frequency aliasing* effect for capturing the feature of $f_{C,N}$. Specifically, frequency aliasing is the phenomenon that allows the data sampler to still retrieve the signal’s spectral feature even when the sampling frequency is lower than the Nyquist frequency (e.g., $2f_{C,N}$). That is, frequency aliasing allows us to assess the frequency even when the frequency is beyond $0.5f_s$. The aliased frequency f^a of the targeted frequency is a series of frequency components. That is, the aliased N^{th} -order harmonics of a C -cylinder engine can be derived as:

$$f_{C,N}^a = |f_{C,N} - K \cdot f_s|, \quad 0 \leq f_{C,N}^a \leq f_s/2 \quad (3)$$

where $K \in \mathbb{Z}$ can be any integer. Furthermore, another key challenge for using the gyroscope sensor is the changing smartphone posture, e.g., mounted in a phone holder or resides in the cup holder. This can induce drastically different gyroscope reading as the raw gyroscope format is a vector of three axis [$Gyro_{roll}$, $Gyro_{pitch}$, $Gyro_{yaw}$]. We calibrate the sensor reading by multiplying the three-dimensional gyroscope reading with a rotation matrix as described in [22]. Now, we can obtain a calibrated one-dimensional gyroscope data which unifies readings from three axes into the fixed yaw axis.

D. Spectral Features of Engine Vibration

By combining Eqs. (1)–(3), $f_{C,N}^a$ of an IC engine at a specific engine speed RPM can be expressed as:

$$f_{C,N}^a(RPM) = \left| \frac{N}{60} RPM \cdot \frac{C}{2} - K \cdot f_s \right|, \quad 0 \leq f_{C,N}^a(\cdot) \leq f_s/2. \quad (4)$$

We validated Eq. (4) by analyzing the gyroscope readings collected from real-world field-tests. We collected data from an

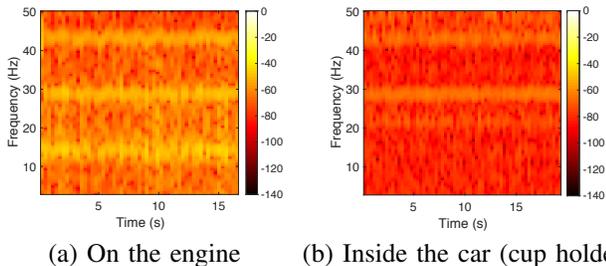


Fig. 3. The idling engine’s spectrogram obtained by gyroscope sensor sampling at 100 Hz with the device residing at different locations.

SUV with a 6-cylinder engine ($C=6$). The sampling frequency f_s of the smartphone’s gyroscope is 100Hz and can thus capture frequencies up to 50 Hz. Since the engine of this idling vehicle (also called *idle speed*) ran at around 880 RPM, the first three harmonics are $f_{6,1}=44$ Hz, $f_{6,2}=88$ Hz, and $f_{6,3}=32$ Hz according to Eq. (2). Based on Eq. (4), the aliased frequencies that can be captured by the smartphone are $f_{6,\{1,2,3\}}^a(880) = \{44, 12, 32\}$ Hz.

Fig. 3 validates $f_{6,\{1,2,3\}}^a(880)$ by showing the spectrogram (frequency components in different time slots) of the field-test gyroscope readings. Specifically, to inspect the engine’s vibration without disturbance induced by other vehicular components (e.g., chassis, suspension, etc.), we opened the engine hood and placed the smartphone directly on the idling vehicle’s engine. Fig. 3(a) shows three major spectral shapes of the engine vibration in 43 Hz, 32 Hz, and 12 Hz frequency bands, which correspond to $f_{6,\{1,2,3\}}^a(880)$. We also emulated the real-world smartphone usage by putting the smartphone in the cup holder inside of the cabin. Although the magnitude of vibration is dimmed (as shown in Fig. 3(b)) due to the propagation from the engine to the vehicle cabin, one can still interpret the same harmonics from gyroscope readings. This observation indicates that the harmonics of an idling vehicle’s engine vibration are stable despite different placements of the sampler (i.e., the smartphone).

E. Classifying Engines

Based on the harmonic feature, we propose an training-less approach that only uses a lookup table — i.e., does *not* require statistical training — for engine characterization and coarse-grained vehicle identification. Note that the cylinder count information can help narrow down the search space of vehicle-types, e.g., the majority of heavy-duty pickup trucks use 8-cylinder engines, while most compact and mid-size cars use 4-cylinder engines.

Building a Lookup Table. We construct a lookup table with spectral features (i.e., extracted by using Eq. (4)) of different engine configurations. Specifically, given the *idle speed* of most vehicles ranging from 600 to 1000 RPM [23], we can derive the *detectable* frequencies (i.e., those that can be captured by the gyroscope at a 100 Hz sampling rate) of the first three harmonics of engines for a different number of cylinders as shown in Fig. 4(a). Note that the fundamental frequency (i.e., the first harmonic) has the strongest magnitude among all other components. We tested this insight by using experimental results as shown in Figs. 4(b), (c), and (d), where we highlight the first harmonic of idling engines with different numbers of cylinders.

Determining the Number of Cylinders. Based on the domain knowledge, an attacker can classify idling engines as follows. In the first step, the attacker needs to extract the idling engine’s

vibration data. The idling period can be extracted when the vehicle is stationary, e.g., in parking lots or waiting for traffic lights at intersections. In the second step, the frequency ranges of different harmonics can be detected with existing spectral analysis algorithms. We tested the performance of linear predictive coding (LPC) [24], a popular spectral analysis in speech recognition, for extracting the frequency of different harmonics. Fig. 4(e) depicts the spectrogram of a Toyota Camry (4-cylinder) engine. We applied LPC on each DFT batch and overlay detection results on the spectrogram. LPC can accurately detect the three harmonics with the averaged central frequency at 24.51 Hz, 43.26 Hz, and 16.65 Hz. In the last step, the attacker can use the lookup table from Fig. 4(a) to match the detected harmonic feature with the engine configuration. The reasoning process can be structured into the decision tree in Fig. 5.

It is worth noting that with a higher gyroscope sampling rate, the frequency ranges overlap less, which can enhance the classification accuracy.

V. CHARACTERIZING MOVING CARS

A. Vibration Analysis

To tackle the complex vibration pattern while the vehicle is moving, we analyze the whole-body vibration based on the source-filter model [25], which has been successfully used for the analysis of a vehicle’s noise, vibration and harness (NVH) [26]. According to the source-filter model, the generation and transmission of mechanical vibrations are described as a two-stage process (Fig. 1). In the first stage, vibrating objects (e.g., engine, tires, and the turbulence induced by wind) excite vibrations, which then start propagating into the mechanical body. The vibration of moving cars can be more intense than that in idling cars. In the second stage, the body structure (vehicle-body) modulates the vibration signal in its transmission path. The modulation process produces structure-borne vibrations at various frequencies. If the frequency of a vibration is close to the natural frequency of the mechanical system, it is called the *resonant frequency* of the mechanical system. Note that the mechanical system vibrates with stronger amplitude [27] at the resonant frequency than other signals.

The resonant frequency is one of the key signatures for characterizing the physical feature of the vibrating object. For example, the resonant frequency of a person’s speech sound has the intrinsic feature of the shape of his/her vocal tract which is used as the unique feature for speaker identification [28]. Likewise, to classify vehicle body-types, the key is to extract the resonant frequencies from the gyroscope readings.

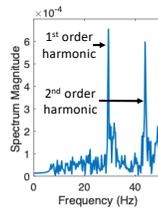
Based on the above findings, we formulate the classification of moving vehicles as a machine-learning problem. As shown in the system pipeline in Fig. 6, the attacker first needs to collect the IMU data traces. This data collection is an one-time effort — the attacker can speed up the process by focusing on the collection of data from the vehicle- types of interest. Next, the data pre-processing and feature extraction module distill the representative features that reflect the vehicle’s vibration pattern, which can subsequently be used for training a statistical model. Next, we present the design of data pre-processing and feature extraction.

B. Data Pre-processing

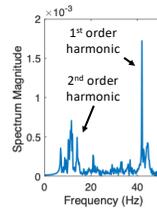
As discussed above, the low-frequency components are noisy and may not be helpful for inferring vehicle-types. The majority

Harmonic order \ # cylinder	N=1	N=2	N=3
4	[20, 33.3]	[33.3, 50]	[0, 40]
6	[30, 50]	[0, 40]	[0, 50]
8	[33.3, 50]	[0, 33.3]	[0, 50]

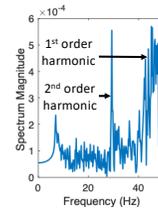
(a) Harmonics of different engine-types



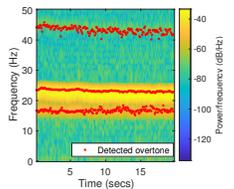
(b) 4-cylinder



(c) 6-cylinder



(d) 8-cylinder



(e) LPC result

Fig. 4. Characterizing the number of engine cylinders based on harmonic features. (a) shows harmonics (derived from Eq. (4)) of different engines. (b), (c), and (d) show the first two harmonics' power spectrums of 4, 6, and 8-cylinder engines, respectively.

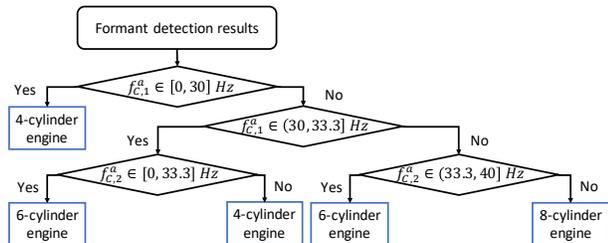


Fig. 5. The *training-free* engine classification based on harmonic features.

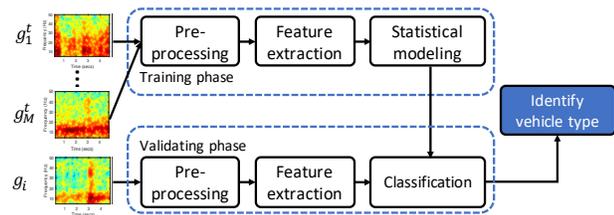


Fig. 6. Pipelined body-type characterization.

of low-frequency vibrations are caused by the driving maneuvers and suspension vibration, for example, due to rough pavements. Driving maneuvers like turns and lane-changes normally last for a few seconds, which create low-frequency signals (≤ 2 Hz) in the gyroscope readings, whereas the suspension vibration caused by rough pavements creates signals of frequency below 5~7 Hz [29]. So, we use a first-order high-pass filter with the cut-off frequency of 7 Hz to remove the low-frequency vibration data.

C. Feature Extraction

We extract the spectral feature of vehicle vibration with a similar workflow inspired by speaker recognition research. The first step is to segment the pre-processed data with a sliding window. Discrete Fourier Transform (DFT) is then used to transform the time-series data to a spectrum distribution (i.e., DFT coefficients in the frequency domain). The spectrum will then be filtered by a series of band-pass filters to extract distinguishable frequency bands. Finally, we calculate the logarithm and Discrete Cosine Transform (DCT) sequentially to derive power spectrum coefficients [28], which form the final feature vector. In what follows, we introduce the design of the data segmentation and filter banks.

1) *Segmentation of Vibration Data*: We slice the gyroscope data trace into short data snippets with a *half-overlapped* sliding window. To extract sufficient data for the analysis in each window, the window length is set to 5 seconds or 500 data samples ($f_s=100$ Hz).

2) *Design of Filter Banks*: Filter banks [30] decouple the spectral data into different frequency bands by processing the signal with a series of bandpass filters. For example, Mel-frequency filter banks (MFCC) [31] have been used to emulate the human ears' perception of sound waves (i.e., air vibration). Here, we use half-overlapped triangular band-pass filters but with an emphasis on the analysis of the vibration signal, which has different frequency range and pattern from the voice signal. To this end, we have designed 15 (i.e., $L=15$) filter banks based on [32]. To meet the requirement of vibration analysis, we tune the frequency span and place triangular filter banks linearly [33] between the lower bound (7 Hz, as shown in Sec. V-B) and the Nyquist frequency. Then, for each filter bank's output, we calculate the summation and take the logarithm of the spectral coefficients. Finally, we use DCT to extract the power spectrum. That is, for each time window, we extract a 1×15 feature vector that reflects the spectral pattern of a data snippet.

D. Classification of Vehicle Types

We apply the thus-derived feature vector on three different machine learning classifiers: Gaussian Mixture Model (GMM), Support Vector Machine (SVM), and Random Forest (RF). In Sec. VII, we will compare the performance of different classifiers. To emulate real-life attack scenarios, we used a *leave-one-out* scheme to construct the training and testing sets. Specifically, for each vehicle type in Table I, we chose one vehicle's data for testing and used the remaining vehicles' for training. To avoid an *unbalanced* dataset, we extracted the same number of data snippets for the training data of each vehicle body-type.

VI. EXPERIMENTAL SETUP

A. Data Collection App

Our data collection emulates real-world smartphone usage while driving. Specifically, our implementation emulates a freemium *driving behavior* analysis app, a fast-emerging mobile app category [34], [35] that actively collects users' motion sensor and location data for analyzing drivers' behavior (e.g., speeding and distracted driving) that is essential for driving-related services and businesses like auto insurance, transportation regulation, etc. For example, when the app is running in the foreground, it may show real-time statistics of the current driving dynamics, e.g., driving behavior detection. Note that to facilitate continuous sensing, the data collection and upload functionalities need to be *always-on*, meaning that unless the user kills the app completely (e.g., terminate the task thread), the app collects the data even when it is running in the

background. To collect the data of moving vehicles, participants were instructed to run the app during their driving.

The key challenge of the app development is controlling the overhead — a high workload on the user device may undermine data collection by exhausting the device’s battery and/or the user’s data plan. Our data-collection app addresses this challenge by regulating the data-collection behavior. In particular, the app samples the gyroscope and accelerometer data at 100 Hz, as shown in Sec. VII-C, a moderate sampling rate that does not incur high energy and CPU overheads. According to our field-test results, the data collection is shown to generate only 4.6 MB data per hour. To lower the user’s cellular data consumption and make the app stealthier, the app uploads the data only when a Wi-Fi connection is available.

Protecting our participants’ privacy is another top priority of our experiment. To achieve this, we applied our university’s IRB and received an approval (Registration No. REDACTED). To ensure our participants are fully aware of the methodology and purpose of our data collection, the app requires each participant’s consent — s/he needs to agree to the terms of use and privacy policy in the consent screen, prior to starting the data-collection process.

B. Collecting Driving Data

The design and validation of our proposed data analytics pipeline (Secs. IV and V) are based on the real-world data collected from field tests. Specifically, we investigated multiple vehicles of different types, and collected data from 17 different cars as listed in Table I. We recruited 10 participants (7 males and 3 females of age ranging from 24 to 58) who use their vehicle for daily commute. Note that the number of vehicles is larger than the number of participants since for some vehicle-types, we also let participants collect data from rental cars. Two types of Android phones (i.e., Google Pixel and LG Nexus 5X) were used for the data collection.

We collected data over 3 months in the U.S. including urban, suburban, and highway environments. The data collection was conducted between 7:30am to 6:30pm. As a result, the accumulated driving dataset has 20.4 hours driving time and 1134.9 km driving distance. The last two columns of Table I denote the travel distance and time for data collection. The average driving speed of compact car, mid-size, SUV, and pickup truck were 67.2, 56.6, 47.8, and 65.6 km/h, respectively.

VII. EXPERIMENTAL RESULTS

A. Differentiating Idling Vehicles

We test the proposed idling vehicles classification using test vehicles of 3 different engine-types. We use eleven 4-cylinder cars (i.e., all compact vehicles and 5 mid-size vehicles), three 6-cylinder cars (i.e., two Ford Explorer SUVs and Mercedes), and three 8-cylinder cars (i.e., all pickup trucks). To emulate real-world smartphone usage in our experiments, smartphones were either mounted on the windshield or placed in the cup holder in each vehicle. Since our current evaluation set is unbalanced, i.e., number of 4-cylinder engine cars is larger than that of 6- and 8-cylinder vehicles, we report precision, recall, F-1, and support to assess the performance of VeFi.

Table II shows the performance of our engine-type classification. Specifically, classification of 6- and 8-cylinder

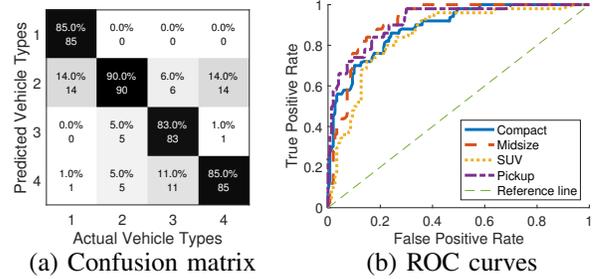


Fig. 7. Confusion matrix and ROC for differentiating vehicle body-types.

engines has zero false positive detection, whereas classification of 4-cylinder engines has zero false negative detection. The main reason for the misclassification of 4-cylinder engines is the large variation of RPM. That is, a larger range of varying RPM may push the first formant closer to the classification boundary.

B. Differentiating Moving Vehicles

To evaluate the performance of moving vehicle characterization in real-world settings, we take a *leave-one-out* approach (as stated in Sec. V) for emulating the attacker’s access of users’ data. That is, for each vehicle body-type, we randomly select one vehicle’s data for testing, while using the remaining vehicles’ data for the statistical modeling. Note that this leave-one-out approach also ensures the training data to be balanced. We repeat this experiment 100 times and report the results in the confusion matrix (see Fig. VII-B(a)). In each iteration, we segment the data and select 500 data snippets for training each vehicle’s vibration profile. With the RF classifier, we can achieve {precision, recall, F-1 score} at {0.85,1.00,0.92} for compact cars, {0.90,0.73,0.80} for mid-size sedans, {0.83,0.93,0.88} for SUVs, and {0.85,0.83,0.84} for pickup trucks, respectively. The overall accuracy is 85.75% based on Fig. VII-B.

To demonstrate VeFi’s performance with a varying threshold, we also present the ROC curve. To plot the ROC curves for multi-class classification, for each targeted vehicle body-type, we represent all other vehicle body-types as *negative* samples. Fig. VII-B (b) shows the resulting ROC curves, where the performance of classifying all vehicle body-types is well above the reference (random guess). The areas-under-curve (AUCs) for compact, mid-size, SUV, and pickup truck ROC curves are 0.8886, 0.9135, 0.8497, and 0.9199, respectively.

1) *Performance of Different Classifiers*: We now compare the performance of different machine-learning classifiers (i.e., RF, SVM, and GMM) under the same experimental settings. We report the average metrics of four different vehicle body-types for each classifier.

As shown in Table III, RF has a similar performance as GMM, and both of these classifiers outperform SVM, possibly because SVM is intrinsically based on distance, which is ill-suited for a spectral analysis.

2) *Top Features for Vehicle Identification*: One of the most useful byproducts of the RF classifier is the measurement of importance [36], [37], which is known to be a good metric of assessing the contributions of different variables to the classification. Specifically, the importance rating of each variable is estimated by using the Gini index [38] — the higher the rating, the more important the corresponding variable’s contribution to the classifier.

Categories	Body Type	Experimental Vehicle(s)	Distance(km)	Hours
C-1	Compact	2009 Toyota Corolla, 2008 Hyundai Elantra, 2018 Nissan Sentra	127.7	1.9
C-2	Mid-size	2006 Honda Accord, 2013 Honda Accord, 2010 Toyota Camry, 2011 Toyota Camry, 2018 Ford Fusion, 2016 Mercedes Benz C-Class	401.2	7.1
C-3	SUV	2013 Honda CRV, 2014 Honda CRV, 2014 Jeep Compass, 2011 Ford Explorer, 2016 Ford Explorer	336.9	7.2
C-4	Pickup truck	2015 GMC Sierra, 2016 GMC Sierra, 2017 Ford F-150	269.1	4.1

TABLE I
DATA COLLECTION FOR DESIGNING AND VALIDATING THE VEHICLE CHARACTERIZATION SCHEME.

# cylinder	Precision	Recall	F-1	Support
4	0.73	1	0.84	11
6	1	0.60	0.75	3
8	1	0.75	0.86	3

TABLE II
PERFORMANCE FOR CLASSIFYING IDLING VEHICLES.

Classifier	Precision	Recall	F-1
RF	0.90	0.86	0.88
GMM	0.89	0.81	0.83
SVM	0.67	0.73	0.72

TABLE III
PERFORMANCE OF DIFFERENT CLASSIFIERS.

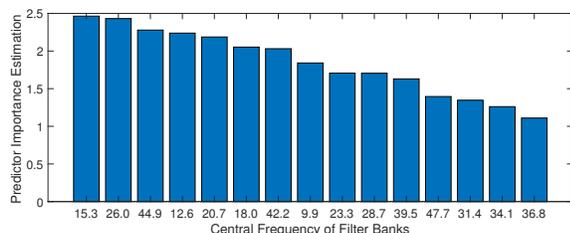


Fig. 8. Importance of variables found by using RF.

In the experiment, we used the central frequency to differentiate filter banks for the feature extraction, where each central frequency corresponds to a variable in the feature vector. We sorted the importance of variables and presented them as a bar plot. Based on Fig. 8, the top-two most important variables are derived from 15.3 Hz and 26 Hz frequency bands. This suggests possible ways of improving the classifier’s performance. For example, one may highlight the important frequency bands by increasing the weight of the corresponding filter banks.

C. Overhead on Smartphones

CPU usage. We recorded the CPU usage on Google Pixel (1.6GHz quad-core CPU) and LG Nexus 5X phones (hexa-core CPU with four 1.4GHz Cortex-A53 and two 1.8GHz Cortex-A57) by using the Android Developer Bridge (ADB) shell. To profile VeFi’s battery overhead, we used Google’s Battery Historian tool [39], which allows developers to inspect the battery usage of each app/module from the bug report generated by smartphones. Note that the average driving time per day in the U.S. is 50.6 minutes. Hence, we generate the bug report from smartphones with 50-min usage of the VeFi app. To evaluate the overhead incurred by sampling and collecting the gyroscope data, we tested the CPU usage and battery drain with smartphones at three different sampling frequencies, i.e., 15 Hz, 100 Hz, and 200 Hz. Here,

Model	Metric	15 Hz	100 Hz	200 Hz
Pixel	CPU usage	1.31%	5.10%	11.20%
	Battery consumed	0.66%	1.30%	2.12%
Nexus 5X	CPU usage	1.36%	6.02%	10.91%
	Battery consumed	0.71%	1.86%	2.80%

TABLE IV
THE AVERAGED CPU AND BATTERY USAGE OF VEFi.

15 Hz IMU sensor sampling rate is the normal rate for many benign smartphone apps (e.g., step counter). 100 Hz is the sampling rate for enabling VeFi. 200 Hz allows more information, which is also the highest sampling rate for many smartphone models.

We report the performance in Table IV. By comparing the metrics between 15 Hz and 100 Hz, the extra CPU and battery usages on Pixel are 3.79% and 0.64%, respectively, whereas Nexus 5X shows 4.66% extra CPU usage and 1.15% more battery consumption. Hence, this marginal increase from 15 Hz to 100 Hz would be hard for the victims to notice. However, increasing the sampling rate from 15 Hz to 200 Hz can drain the victims’ smartphone battery much faster, thus making it difficult for the attacker to stealthily collect data.

VIII. RELATED WORKS

Privacy Breach via Mobile Sensory Data. Michalvesky *et al.* [40] proposed Gyrophone, a scheme for identifying speakers and reconstructing simple spoken words by harvesting the gyroscope reading close to the speaker. Roy *et al.* [41] proposed VibraPhone, a series of data analytics techniques for reconstructing photonic information from vibration data induced by vibrating motors in mobile devices. In essence, the authors analyzed the changes in IMU readings that are induced by external disturbances (i.e., sound waves). Dey *et al.* [42] showed that the imperfection of the MEMS chip is distinguishable among mobile devices. Narain *et al.* [43] analyzed the location-privacy threat via leakage of smartphone IMU data. To the best of our knowledge, there does not exist any work that investigates the feasibility of analyzing the vehicle’s mechanical vibration and its connection with the vehicle’s physical features. **Vehicle Vibration Analysis.** Sun [44] introduced the sensor fusion of vehicle vibration signals and oil data to monitor vehicle health. Puchalski [45] investigated a series of statistical models of using vibration signal for vehicle diagnostics. Komorska [46] introduced a scheme of using the resonant frequency for detecting mechanical defects of a vehicle and/or an engine. Kozhisseri *et al.* [47] proposed an acoustic feature for identifying a vehicle. Specifically, the key feature they extracted is the fundamental frequency of idling engine vibrations. Goksu *et al.* [48] analyzed vibro-acoustic data with wavelet decomposition

to characterize a vehicle engine while the vehicle is moving. Unlike these prior works, we explore the feasibility of extracting rich physical features from vehicles based on smartphone IMUs.

IX. CONCLUSION

We have presented VeFi, a novel attack that can characterize vehicles by only using motion sensor data generated by driving apps on users' smartphones. VeFi only uses zero-permission IMU sensor readings and incurs minimal resource overhead to the mobile device, enabling a stealthy side-channel attack with serious privacy implications.

ACKNOWLEDGMENTS

This work was supported in part by the Chinese NSFC Grant No. 62102256 and the US ONR under Grant No. N00014-22-1-2622.

REFERENCES

- [1] Do You Drive a Liberal Car or Conservative Car? <https://baristanet.com/2012/08/do-you-drive-a-liberal-car-or-conservative-car/>.
- [2] Sangho Choo and Patricia L Mokhtarian. What type of vehicle do people drive? the role of attitude and lifestyle in influencing vehicle type choice. *Transportation Research Part A: Policy and Practice*, 38(3):201 – 222, 2004.
- [3] What your car says about you. <https://www.forbes.com/2009/10/06/car-personality-wealth-lifestyle-vehicles-gender-income.html#1d301337399b>.
- [4] How Cambridge Analytica turned Facebook 'likes' into a lucrative political tool. <https://www.theguardian.com/technology/2018/mar/17/facebook-cambridge-analytica-kogan-data-algorithm>.
- [5] Nyquist-Shannon Sampling Theorem. https://en.wikipedia.org/wiki/Nyquist%E2%80%93Shannon_sampling_theorem.
- [6] Sensor Penetration in Smartphones Shipped in 2017. <https://www.counterpointresearch.com/sensors-smartphones-top-10-billion-unit-shipments-2020/>.
- [7] Coriolis force. https://en.wikipedia.org/wiki/Coriolis_force.
- [8] Ville Kaajakari et al. Practical mems: Design of microsystems, accelerometers, gyroscopes, rf mems, optical mems, and microfluidic systems. *Las Vegas, NV: Small Gear Publishing*, 2009.
- [9] Here's What Really 'Drives' Democrats And Republicans. <https://www.forbes.com/sites/jimgorzelay/2016/11/02/heres-what-really-drives-democrats-and-republicans/#19a7e85e22b8>.
- [10] Moscow Billboard Targets Ads Based on the Car You're Driving. <https://www.technologyreview.com/s/603743/moscow-billboard-targets-ads-based-on-the-car-youre-driving/>.
- [11] Dena M. Gromet, Howard Kunreuther, and Richard P. Larrick. Political ideology affects energy-efficiency attitudes and choices. *Proceedings of the National Academy of Sciences*, 110(23):9314–9319, 2013.
- [12] Timnit Gebru, Jonathan Krause, Yilun Wang, Duyun Chen, Jia Deng, Erez Lieberman Aiden, and Li Fei-Fei. Using deep learning and google street view to estimate the demographic makeup of neighborhoods across the united states. *Proceedings of the National Academy of Sciences*, 114(50):13108–13113, 2017.
- [13] Your cars: politics on wheels. <https://www.nytimes.com/2005/04/01/automobiles/your-car-politics-on-wheels.html>.
- [14] Gallup Poll Analysis: Political Correlates of Car Choice. <http://news.gallup.com/poll/23230/gallup-poll-analysis-political-correlates-car-choice.aspx>.
- [15] The New York Times: How Trump Consultants Exploited the Facebook Data of Millions. <https://www.nytimes.com/2018/03/17/us/politics/cambridge-analytica-trump-campaign.html>, 2018.
- [16] US Department of State: Does Everyone in America Own a Car? https://photos.state.gov/libraries/cambodia/30486/Publications/everyone_in_america_own_a_car.pdf.
- [17] U.s. energy information administration - eia - independent statistics and analysis.
- [18] Tutorial of Engine Power Train. <https://www.micksrgarage.com/blog/how-spark-plugs-work>.
- [19] V Ganesan. *Internal combustion engines*. McGraw Hill Education (India) Pvt Ltd, 2012.
- [20] Richard Feynman. *The Feynman Lectures on Physics: Volume 2*, volume 2 of *The Feynman Lectures on Physics*. Addison-Wesley, Boston, 1963.
- [21] J. Yang, Y. Suematsu, and Z. Kang. Two-degree-of-freedom controller to reduce the vibration of vehicle engine-body system. *IEEE Transactions on Control Systems Technology*, 9(2):295–304, March 2001.
- [22] Dongyao Chen, Kyong-Tak Cho, Sihui Han, Zhizhuo Jin, and Kang G. Shin. Invisible sensing of vehicle steering with smartphones. In *Proceedings of the 13th Annual International Conference on Mobile Systems, Applications, and Services*, MobiSys '15, pages 1–13. ACM, 2015.
- [23] Idle speed. https://en.wikipedia.org/wiki/Idle_speed, 2016.
- [24] D. O'Shaughnessy. Linear predictive coding. *IEEE Potentials*, 7(1):29–32, Feb 1988.
- [25] J. Durrieu, G. Richard, B. David, and C. Fevotte. Source/filter model for unsupervised main melody extraction from polyphonic audio signals. *IEEE Transactions on Audio, Speech, and Language Processing*, 18(3):564–575, March 2010.
- [26] Maria Antonietta Panza. A review of experimental techniques for nvh analysis on a commercial vehicle. *Energy Procedia*, 82:1017 – 1023, 2015. 70th Conference of the Italian Thermal Machines Engineering Association, ATI2015.
- [27] G.S. PADDAN and M.J. GRIFFIN. Evaluation of whole-body vibration in vehicles. *Journal of Sound and Vibration*, 253(1):195 – 213, 2002.
- [28] Douglas Reynolds and Richard Rose. Robust Text-Independent Speaker Identification Using Gaussian Mixture Speaker Models. In *IEEE Transactions on Speech and Audio Processing*, Jan. 1995.
- [29] Spring rate and suspension frequencies. <https://www.drtoned.com/tech-ramblings/2017/10/2/spring-rates-suspension-frequencies>, 2016.
- [30] M. Vetterli and C. Herley. Wavelets and filter banks: theory and design. *IEEE Transactions on Signal Processing*, 40(9):2207–2232, Sep. 1992.
- [31] G. Tzanetakis and P. Cook. Musical genre classification of audio signals. *IEEE Transactions on Speech and Audio Processing*, 10(5):293–302, July 2002.
- [32] Todor Ganchev, Nikos Fakotakis, and George Kokkinakis. Comparative evaluation of various mfcc implementations on the speaker verification task. In *Proceedings of the SPECOM*, pages 191–194, 2005.
- [33] Tomi H Kinnunen. *Optimizing spectral feature based text-independent speaker recognition*. University of Joensuu, 2005.
- [34] TrueMotion. <https://gotruemotion.com/>, 2017.
- [35] CMT. <https://www.cmtlematics.com/>, 2018.
- [36] Carolin Strobl, Anne-Laure Boulesteix, Achim Zeileis, and Torsten Hothorn. Bias in random forest variable importance measures: Illustrations, sources and a solution. *BMC Bioinformatics*, 8(1):25, Jan 2007.
- [37] Kellie J. Archer and Ryan V. Kimes. Empirical characterization of random forest variable importance measures. *Computational Statistics and Data Analysis*, 52(4):2249 – 2260, 2008.
- [38] M. Pal. Random forest classifier for remote sensing classification. *International Journal of Remote Sensing*, 26(1):217–222, 2005.
- [39] Inc. Google. Battery historian. <https://github.com/google/battery-historian>, 2017.
- [40] Yan Michalevsky, Dan Boneh, and Gabi Nakibly. Gyrophone: Recognizing speech from gyroscope signals. In *Proceedings of the 23rd USENIX Conference on Security Symposium, SEC'14*, pages 1053–1067, Berkeley, CA, USA, 2014. USENIX Association.
- [41] Nirupam Roy and Romit Roy Choudhury. Listening through a vibration motor. In *Proceedings of the 14th Annual International Conference on Mobile Systems, Applications, and Services*, MobiSys '16, pages 57–69. ACM, 2016.
- [42] Sanorita Dey, Nirupam Roy, Wenyuan Xu, Romit Roy Choudhury, and Srihari Nelakuditi. Accelprint: Imperfections of accelerometers make smartphones trackable. In *NDSS*, 2014.
- [43] Sashank Narain, Triet D. Vo-Huu, Kenneth Block, and Guevara Noubir. The perils of user tracking using zero-permission mobile apps. *IEEE Security Privacy*, 15(2):32–41, 2017.
- [44] Qiao Sun. Sensor fusion for vehicle health monitoring and degradation detection. In *Proceedings of the Fifth International Conference on Information Fusion. FUSION 2002. (IEEE Cat.No.02EX5997)*, volume 2, pages 1422–1427 vol.2, July 2002.
- [45] Andrzej Puchalski. A technique for the vibration signal analysis in vehicle diagnostics. *Mechanical Systems and Signal Processing*, 56-57:173 – 180, 2015.
- [46] I. Komorska. Utilising the resonance frequency of the engine vibration sensor in diagnostics of the exhaust valve leakage. *Journal of KONES*, Vol. 17, No. 2:209–216, 2010.
- [47] S. Kozhisseri and M. Bikkdash. Spectral features for the classification of civilian vehicles using acoustic sensors. In *2009 IEEE Workshop on Computational Intelligence in Vehicles and Vehicular Systems*, pages 93–100, March 2009.
- [48] Huseyin Goksu. Engine speed-independent acoustic signature for vehicles. *Measurement and Control*, 51(3-4):94–103, 2018.