Driver Identification via Brake Pedal Signals — A Replication and Advancement of Existing Techniques

Bernhard Gahr[1], Benjamin Ryder[2], Andre Dahlinger[1] and Felix Wortmann[1]

***Abstract*— Driver identification is a growing topic which offers a streamlined user experience in the connected car, but potentially also highlights privacy issues of our interconnected lives. Recent studies have reported the ability for individuals to be reliably identified out of a group based on their driving behavior. In particular, the state-of-the-art study claims that, in a controlled setting, data collected on how a driver operated the brake pedal could perfectly distinguish between 15 drivers. The paper at hand was not able to validate these strong scientific claims using naturalistic driving data. In line with the results of other studies using similar data, the replicated identification accuracy dropped to values between 40% and 70% by applying the outlined methods. Nevertheless, this paper further contributes by adapting the reported feature collection technique in order to achieve identification results between 80% and 99.5% in this challenging setting, thus advancing the state-of-the-art. These findings demonstrate the real-world capabilities of data-enabled driver identification, which both facilitates new use-cases and potentially raises privacy questions. As such, important key features from the identification models are presented to assist both researchers and practitioners in this rapidly developing topic.**

I. Introduction

The recent work of multiple research groups shows the growing interest of driver identification. A seamless driver identification fuels the hope of a streamlined driver experience without the need of a cumbersome authentication, and hence could be used to increase security as an intrusion detection system, or enable personalized service models [1]. At the same time an automated identification raises questions regarding privacy concerns, especially because of the increasing interconnectivity of modern cars. As in other fields, such as credit card payments [2], health-related data [3] or movie ratings [4], recent studies showed the high potential of identifying individuals reliable out of a group using car data. In order to leverage the positive aspects and prevent the risks of person identification, it is important to understand its technical background and potential. This paper contributes by investigating person identification in the realm of the car, through data collected from the CAN bus of a fleet of 50 vehicles in a naturalistic setting.

Lab based simulation studies, where users are monitored in a driving simulator, and controlled settings, where drivers follow a set route and perform specific maneuvers, both provide valuable insights into driving behavior. However, in practice it is important that results are validated in more realistic and challenging real-world driving situations, where drivers operate freely and under no instruction, otherwise known as naturalistic driving field studies.

One of the most recent studies, where brake pedal data was collected in a controlled setting, claims perfect identification rates when distinguishing between 15 drivers [5]. As such, the validation of this strong scientific claim is an important academic task, and the paper at hand was not able to replicate the reported results. By using the state-of-the-art methods outlined, and naturalistic driving data, the identification accuracy decreased to performances indicated in other studies using similar data.

However, this does not imply that a strong identification using only the brake pedal signal is impossible. Therefore, this paper further investigates an alternative feature collection method, motivated by an information theoretic description of the recorded signals due to the typically sparse usage of the brake pedal over time. The adapted method achieves higher identification rates, and thus advances the state-of-the-art.

The paper is structured as follows: First, we present related research on the topic of car data-based person identification. Then we describe the study design and our field test settings, followed by the analytical methodologies we apply. Further, we present our results validating existing methods under a naturalistic setting and showing results applying the novel feature collection method. In the final section, we discuss the results and shed some light on the implications of this research.

II. Related Work

The majority of the related research on driver identification in the recent years is motivated by the beneficial potential for improving driver safety, rather than raising awareness of potential privacy issues. Similar to the approaches explored in various other fields outside of the automotive industry, researchers have tried to identify drivers using various different methods. For example, CAN bus data has been used to classify actions, detect distraction and identify drivers using their driving-behavior characteristics [6]. A deeper analysis of driver re-identification using a simulator led to the identification of drivers with an accuracy of up to 73% [7]. To achieve this result, the authors consider driving behavior signals, such as the accelerator pedal, brake pedal, vehicle velocity and following distance. Later, the same group was able to optimize re-identification to 89.6% for the simulator, which dropped to 71% for real-car data, by applying spectral analysis methods [8]. In a study of 276 drivers measuring

*This work was supported by Robert Bosch GmbH

[1]Institute of Technology Management, University St. Gallen, Switzerland {bernhard.gahr, andre.dahlinger, felix.wortmann}@unisg.ch

[2]Department of Management, Technology and Economics, ETH Zurich, Switzerland bryder@ethz.ch

simulator- and real-car signals as well as the following distance, the group was able to re-identify a driver with 76.8% accuracy [9]. And further, a recent study focused on the technical implementations of driver identifying methods [10]. The drivers in this lab study had to follow a predefined route of 25 km. Results were limited to detecting subsets of the available 11 drivers, reaching identification performances from 52% to 98% for different settings. To achieve this, they used an extreme learning machine network which was applied on audio-, video-, inertial measurement unit-, frontal laser scanner-, and CAN bus signals. Other approaches to identify drivers used the sitting posture [11], the voice of the driver in a vehicle security system [12], and finger-vein technology by applying a neural network [13].

However, for privacy reasons it is important to explore the potential of driver identification using solely CAN bus signals. As such, one recent study was exploring driver identification using only CAN bus signals, such as steering wheel angle, steering velocity, throttle position, and brake pedal position under a naturalistic driving scenario [14]. The authors limited themselves to identify a driver only by looking at the signals generated during a turning event. The data was collected by 64 drivers accumulating 2,098 hours and 110,023 km of driving data. By using a random-forest classifier a subset of 2 to 5 drivers were identified with accuracies between 50% and 77%.

Finally, two most recent studies highlighted the high potential of using CAN bus data for violation of drivers' privacy. In the first study the researchers aimed to identify up to ten drivers in the shortest time frame possible, and achieved perfect identification within two minutes [15]. The study used three different datasets, all containing driving data from predefined routes. In the second study, the authors were able to perfectly identify every participant out of a sample of 15 drivers even when only the brake pedal signal was used, as the drivers followed predefined routes and performed maneuvers on a parking lot [5]. Furthermore, they came to the result that a 99% identification accuracy could be achieved with only 13.5 minutes of driving data for training and 1.5 minutes for testing with the top 5 sensors. Both studies concluded, that the brake pedal signal was among the best indicators for identifying the drivers.

Summing up, we see that studies using data from real-world scenarios, lead to identification rates between 50% and below 80% for 2 to 5 drivers. High identification rates were primarily achieved in lab settings, even when the brake pedal signal was solely used. Since there is such a strong performance difference between results of lab studies and results based on naturalistic data, we focus in a first step on the validation of the lab setting results. Thereby the insights from previous work is used, i.e. the limitation on the brake pedal only and the potential performance improvement when limiting the identification to specific driving moments.

III. Data Collection & Analysis Methods

To validate the results, we conducted a large naturalistic field study. In this field study, we collected CAN bus data from 50 professional road assistance drivers over a period of three months, covering approximately 300,000 kilometers of naturalistic driving. The data collection system used for this study was introduced to the drivers as an additional driving aid, giving feedback on their fuel consumption, which they could use voluntarily. The drivers were not asked to perform any special tasks or take any predefined routes and used the system during their normal daily work. The car's CAN bus was accessed via the OBD-II interface, using a dongle that sent the information via Bluetooth to a smartphone. All data was then transmitted to the server via the cellular network. Among other signals, we collected brake pedal position, measured in percentages (0% = not pushed; 100% = fully pushed) with a maximum frequency of 30 Hz. To validate the appraoch of [5], the signal data was further pre-processed by a quadratic interpolation and resampled at 30 Hz for a constant sampling frequency. In the following two sections we will outline the two fundamental approaches which we applied, in order to validate the state-of-the-art methods, and advance and improve current methods, respectively.

The first approach will be called sliding window feature collection-, the second brake event feature collection approach. In the following vectors denoting sequences of the sliding window approach will be denoted by $\vec{W}_i$ and vectors denoting sequences of the brake detection approach by $\vec{B}_i$.

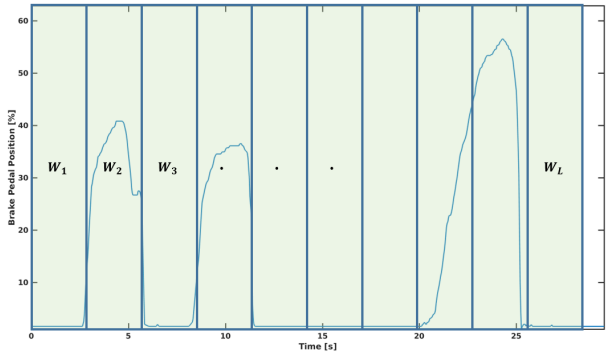


Fig. 1. Schematic of the sliding window approach applied on an incoming brake pedal signal. Each window represents one feature vector.

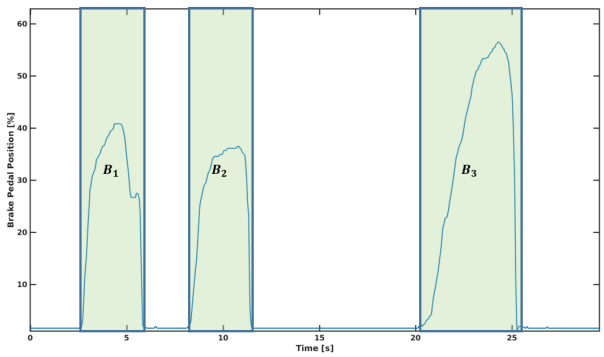


Fig. 2. Schematic of the feature vector detection per brake event. Each braking represents one feature vector.

A. Sliding Window Feature Collection

Based on the work of Enev et al., we replicate the sliding window analysis on the brake pedal signal of the CAN bus [5]. Let a recorded brake pedal signal be denoted by the vector $\vec{X}$ of length N. Hence the i-th element of this vector is denoted by $\vec{X}[i]$ with $i \in [1,N]$, and a subsequence of length n starting at i as $\vec{X}[i, i+n-1]$. The sliding window is applied to brake pedal signals as depicted in Figure 1. The signal in each window is used to collect several features, which are stacked in one feature vector. Hence, each window represents one feature vector. The length of a window was optimized to 3 seconds and an overlapping of the windows of 25%, i.e. 0.75 seconds. The window overlap is denoted by ol. With the specified 30 Hz this leads to a window length (subsequence) of $n = 90$. As such, the sliding-window vector can be denoted as shown in Equation 1 below:

$$\vec{W}_i = \vec{X}[(n-ol)\cdot(i-1), (n-ol)\cdot i + ol - 1] \quad \forall\, i \in \left[1, \left\lfloor \frac{N}{n-ol} \right\rfloor\right] \tag{1}$$

B. Brake Event Feature Collection

The second approach goes beyond the prior work and collects feature vectors based on brake events. A similar approach was used by Hallac et al., where the recorded signals were restricted to turning moments [14]. The authors showed that turns are particularly well-suited for detecting variations across drivers. Similarly, we use braking events in order to distinguish different behaviors among the drivers, and limit the feature collection only to the moments where the brake pedal is applied. Figure 2 shows an illustrative schema of the brake event method. Each brake event, from the moment the brake is applied until the moment it is released, is used to calculate one feature vector. Due to the possible loosening of brake pedal springs, we set a threshold of $th = 8\%$ to trigger a brake event. Hence, $i_{start} = \min\{j \in [1,N] \mid \vec{X}[j] \geq th\}$. The index of when the brake is released can be found by $i_{end} = \min\{j \in [i_{start}+1, N] \mid \vec{X}[j] < th\}$ This leads to the vector of the subsequence of $\vec{B} = \vec{X}[i_{start}, i_{end}]$. To collect all brake events, this can be written as shown in Algorithm 1. Note that, when compared to $\vec{W}_i$ from (1), the vectors $\vec{B}_i$ can cover subsequences of different lengths, as it can be seen in Figure 2 comparing the width of $\vec{B}_1$, $\vec{B}_2$, and $\vec{B}_3$.

Algorithm 1 Brake Event Detection Algorithm

```
Input: X⃗
i = 1
while true do
    N = len(X⃗)
    i_start = min{j ∈ [1,N] | X⃗[j] ≥ th}
    if i_start = ∅ then
        break
    i_end = min{j ∈ [i_start, N] | X⃗[j] < th}
    if i_end = ∅ then
        break
    B⃗_i = X⃗[i_start, i_end]
    X⃗ = X⃗[i_end + 1, N]
    i = i + 1
```

We motivate feature collection by brake events by considering a sliding window where the pedal is not pushed, e.g. $\vec{W}_i = \vec{0}$. Since the probability of the signal being zero is one, the entropy of the signal becomes zero. Following the description of the mutual information [16], it can be shown that $I(D_j, \vec{W}_i) = 0 \;\forall\; j$, where D_j describes the j-th driver. Hence, we have no knowledge gain about the drivers looking at these windows. Similarly, this can be shown for the brake event approach, for subsequences, when the brake pedal is not applied. The only information we exclude using this approach is the time between two brake events. We argue that the duration between two events reveals only information about the surrounding situation, rather than the driver itself, i.e. we assume that the reason *why* a brake was performed is mainly due external factors, but only *how* a brake is performed reveals information about the driver.

The benefit of this approach is further demonstrated when considering the sparsity of brake events over time. As part of our analysis, 1000 brake events were extracted for each of the 50 drivers. The mean time between two brake events lies around 25 seconds, and the maximal amount close to 100 seconds, as shown in Figure 3. Additionally, there were several outliers which were excluded from the figure, with extremely high durations between brake events. The time between these outliers rose to over 20 minutes on trips including highway segments. With the observation of sparse brake events over time, and the property that the time where the brake is not applied does not reveal any information about the driver, we "condensed" the information of a driver into fewer feature vectors. This new approach effectively removes unimportant, and potentially uncertainty increasing, feature vectors to improve classification accuracy.

C. Feature Calculation

In the following we will describe the calculation of the features. In total, 58 features were collected for both approaches. The statistical features contain the signals' minimum (only for sliding-window), maximum, average, quartiles, standard

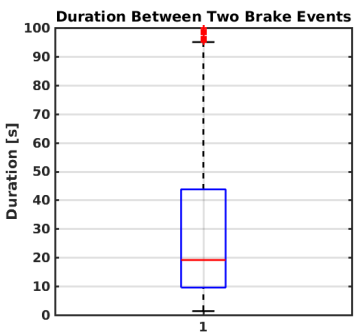


Fig. 3. Distribution of the duration between two brake events.

deviation, autocorrelation, kurtosis, skewness and duration (only for brake event). The descriptive features contain a ten piece-wise approximation of the signal. And the frequency features contain the frequency power components after a Fourier transformation. We write the set of functions to calculate the features as $\vec{x}_i \mapsto f_j(\vec{x}_i)$, where $j \in [1,58]$ (i.e. one of the 58 functions). The vector $\vec{x}_i$ denotes hereby either one of the sliding window vectors $\vec{W}_i$ or one of the brake event vectors $\vec{B}_i$. All functions are mapping the signal vector onto a real number, i.e. $f_j(\cdot) : \mathbb{R}_+^{n\times 1} \to \mathbb{R}$. The features are then stacked into one feature vector, i.e. $\vec{F}_i = [f_1(\vec{x}_i), f_2(\vec{x}_i), ..., f_{58}(\vec{x}_i)]^T$.

D. Classification

Classification of the drivers was performed using various machine learning techniques, including support-vector-machines, k-nearest-neighbors and naive Bayes algorithms. The identification rates of the random-forest algorithm were highest, since it performs well with a larger number of features [17]. As the focus of this paper is on the feature collection method, we exclude the results of the other algorithms in the later results section for brevity, and focus on the random-forest algorithm. We used the *randomforest-matlab* [18] library for *MATLAB* to identify the drivers. The random-forest algorithm constructs multiple decision trees during the training phase. The models can be build in two ways: pairwise, i.e. one model per class [19], or one model for all classes. Despite each model in the pairwise classification being less complex, it was more time-consuming in training and testing, and the results did not significantly increase. Therefore, the results presented use one model with 5,000 trees, where each decision tree returned one vote for which class the input feature vector fits best. The final driver classification decision is made by majority vote. Evaluating the performance with a sequence of N feature vectors per class leads to a majority vote over N*5,000 feature vectors.

IV. Results

The results shown below are structured in the following way: in Section IV-A, the results validating the results from a lab study setting are presented using 5 and 15 drivers. Second, in Section IV-B the performances achieved by the newly introduced brake event feature collection approach are given for groups of 5, 15 and 50 drivers, using the same performance indicators of [14], [5] and leveraging the full dataset. Further the confusion matrix for the identification of 50 drivers is given in Section IV-C. And finally a short description on the feature importance is given, hinting which features are strong indicators of a driver, in Section IV-D.

For each categorization task, the results are averaged over multiple iterations where drivers in the group were picked uniformly at random out of the 50 available drivers from the field study. We also assessed the length of driving time needed to train accurate identification models by using different amounts of feature vectors from the training set in each iteration. Moreover, the amount of test set feature vectors for each iteration was measured, giving an indication of driving time needed until accuracy converged on driver identity. Finally, the feature vectors for each iteration were picked uniformly at random from the training and test set. Naturally, the training and test sets were not overlapping so that the data used to validate the results were unknown to the trained model.

A. Validation of state-of-the-art methods

The replication results of identifying 5 (blue curves), and 15 (red curves) drivers based on the sliding window approach are shown in Figure 4. The dashed curves denote the results for a trained model using 15 minutes of driving data and the solid curves denote the identification accuracy of the model trained with 90 minutes of driving data. We observe a small increase, of less than 5%, in accuracy by using more driving data to train the model, however, we see that this is not a dramatic improvement. Thus, the identification accuracy rises from 20% (random guess) to 70% for 5 drivers and from 6.7% to over 40% for 15 drivers. Further we observe that results using driving data from a lab setting could not be replicated with our dataset [5], but ratifies the reported results of other studies using naturalistic driving data and applying state-of-the-art methods [14].

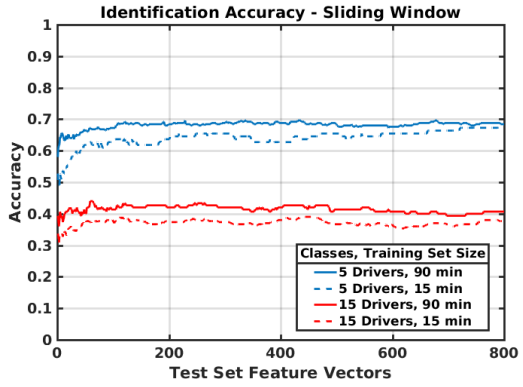


Fig. 4. Identification accuracy over feature vectors collected by the sliding window approach for 5 and 15 drivers.

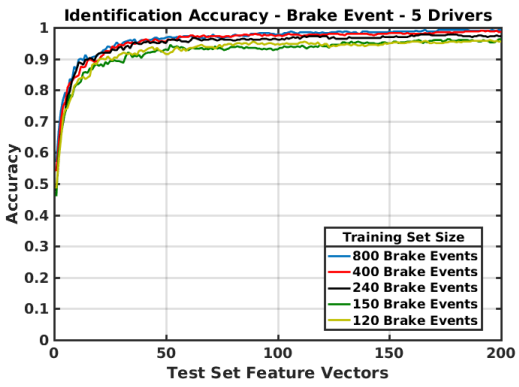


Fig. 5. Identification accuracy over feature vectors collected by the brake event detection approach for 5 drivers.

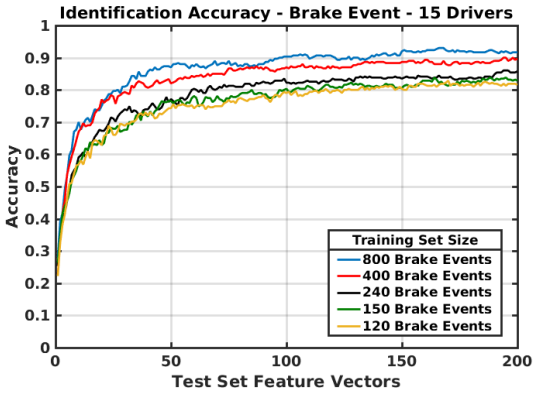


Fig. 6. Identification accuracy over feature vectors collected by the brake event detection approach for 15 drivers.

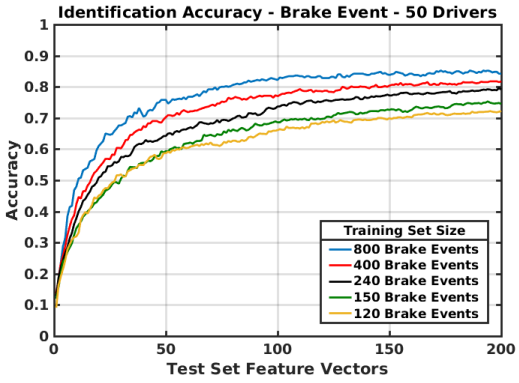


Fig. 8. Identification accuracy over feature vectors collected by the brake event detection approach for 50 drivers.

B. Brake event feature collection method

In the case of the brake event approach, each feature vector corresponds to one brake event, where approximately 80 brake events equates to 30 minutes of driving time. The results of the newly introduced brake event approach are given in Figures 5, 6 and 8. The 200 test set feature vectors shown in these Figures correspond to a driving time of approximately 75 minutes on average. For all of the brake event classification tasks, we present results with a testing set size with a maximum of 200 brake event feature vectors, since the accuracy converges only then for 50 drivers as shown in Figure 8. The training set sizes range from 120 to 800 brake events, corresponding to an approximate range from 45 minutes to 5 hours of driving time according to the statistics shown in Figure 3. We chose a maximum training set size of 800 feature vectors, since this would correspond to an expected driving time of over 5 hours, therefore larger training set sizes are unrealistic.

Figure 5 shows the performance of the random-forest algorithm using the approach of brake event feature collection for 5 drivers. Altogether, the performance increases to over 95% for a training set size of 120 feature vectors, and to 99.25% for 800 feature vectors after using 200 feature vectors from the test set. In comparison to the sliding window approach, we see a significant improvement in identification accuracy. A similar improvement is shown for 15 drivers in Figure 6. The accuracy increases from 40% with the sliding window approach, to between 80% and 93% depending on the training set size. Finally, Figure 8 shows the performance for the full dataset of 50 drivers, where the base line accuracy of randomly guessing the driver identity lies at 2%. From this we see an increase to an accuracy between 70% when training the model with 120 brake events and 85% when using 800 brake events.

Summing up, we observe that adapting the feature collection approach to an event based method, we were able to almost perfectly identify 5 drivers. Naturally, the accuracy of these predictions drop when more drivers are added, and with 45 minutes of training driving data we could identify 15 drivers with an accuracy of 80%, and 50 drivers with an accuracy of 70%.

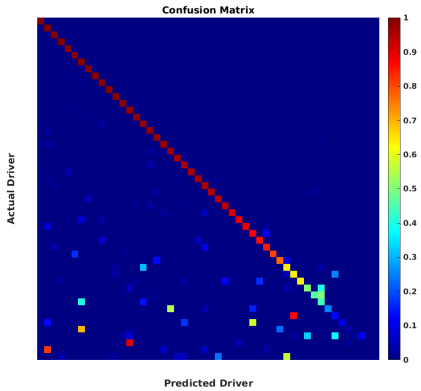


Fig. 7. Confusion matrix of 50 drivers after training with 800 and testing with 200 brake events.

C. Confusion Matrix

Figure 7 shows the confusion matrix of 50 drivers in a sorted order after training the algorithm with 800 brake events and applying 200 brake events for testing. We see that for more than 30 of the drivers, brake events led to a good prediction of the actual driver. Approximately 6 drivers (orange to green squares on the diagonal) were correctly identified between 50% to 80% of the time, and the last 10 drivers were almost never correctly identified. Moreover, it can be seen that there is no single driver which was used as a default prediction, or 'sink', for all misclassified drivers. Three of the poorly identified drivers were each heavily misclassified towards one specific driver, shown by the red squares not on the diagonal. These drivers were mismatched with a driver from a different base location, indicating that there was no over-fitting to the region of operation, but rather that the drivers have a very similar driving style.

D. Feature Importance

Finally, the importance of the features calculated by the random-forest algorithm changed for different selections of

drivers and different sizes of training sets. However for all iterations a few features always appeared to have high importance, such as the first, second, 9th, and 10th piecewise approximations. These features correspond to the downsampled signal at the beginning and the end of a brake event. Additionally, the frequency components between 3 Hz and 8 Hz were strong indicators of the driver.

V. Discussion and Outlook

This work contributes to the growing field of automatized person identification. Applying this topic to the automotive setting could enable and improve many products and services that require driver identification, such as automatic activation of personal insurance plans, frictionless personalization of the car experience, or car intrusion detection systems. Meanwhile, the results presented enable policy makers and companies to better evaluate the potential privacy concerns of connected cars.

Recent reported results from various groups, showed the ability for individuals to be reliably identified out of a group based on their driving behavior. In particular, one of the most recent studies claims that, in a controlled setting, data collected on how a driver operated the brake pedal could distinguish between 15 drivers with perfect accuracy. The validation of such a strong scientific claim is an important aspect of any academic discipline, and the paper at hand was not able to replicate the reported results, but rather verified results using state-of-the-art methods of other studies using naturalistic driving data.

The results of this research should be assessed in light of its limitations. Despite the uniqueness of the data collected, we have to acknowledge that our analysis is based on a field study of professional road assistance drivers. This implies that our results may not be easily generalizable to the wider population. However, the homogeneity of the drivers may have reduced the variance in the feature distribution and actually weakened the overall performance.

Nevertheless, this paper validates the high potential of driver identification using the brake pedal, by adapting the reported feature collection technique in order to achieve higher identification results in this challenging setting, and thus advancing the state-of-the-art. We motivated our approach with an information theoretic learning description, and showed that we can increase the identification performance significantly by only considering the times where the brake pedal was applied. As such, the feature collection approach derived in this paper is tailored to the sparsity of brake events over time. Therefore, it is possible that for each variety of CAN bus signals a different feature collection approach will be more appropriate and hence, we believe that immediate and perfect driver identification, even for a larger set of drivers, can be achieved through combining multiple signals and feature collection approaches.

These findings demonstrate the real-world capabilities of data enabled driver identification, which both facilitates new use-cases and potentially raises privacy questions. As such, we are excited for the results of future work, which should apply the approach presented in this paper to a field setting with a variety of signals.

References

[1] L. F. Wurster, J. M. Correia, A. Woodyer, and C. Eschinger, "Market Trends: Prepare Now to Smooth the Inevitable Transition to a Subscription-Based Business Model." Gartner, Tech. Rep., 2015.

[2] Y.-A. de Montjoye, L. Radaelli, V. K. Singh, and A. S. Pentland, "Unique in the shopping mall: On the reidentifiability of credit card metadata," *Science*, vol. 347, pp. 536–539, 2015.

[3] a. C. Solomon, R. Hill, E. Janssen, S. a. Sanders, and J. R. Heiman, "Uniqueness and how it impacts privacy in health-related social science datasets," *Proceedings of the 2nd ACM SIGHIT symposium on International health informatics - IHI '12*, p. 523, 2012.

[4] A. Narayanan and V. Shmatikov, "Robust De-anonymization of Large Sparse Datasets," in *Proceedings of the 2008 IEEE Symposium on Security and Privacy*, ser. SP '08. Washington, DC, USA: IEEE Computer Society, 2008, pp. 111–125.

[5] M. Enev, a. Takakuwa, K. Koscher, and T. Kohno, "Automobile Driver Fingerprinting," *Proceedings on Privacy Enhancing Technologies*, vol. 2016, no. 1, pp. 34–50, 2016.

[6] S. Choi, J. Kim, D. Kwak, P. Angkititrakul, and J. H. L. Hansen, "Analysis and classification of driver behavior using in-vehicle can-bus information," *Biennial Workshop on DSP for In-Vehicle and Mobile Systems*, no. October 2015, pp. 17–19, 2007.

[7] T. Wakita, K. Ozawa, C. Miyajima, K. Igarashi, K. Itou, K. Takeda, and F. Itakura, "Driver identification using driving behavior signals," in *IEEE Conference on Intelligent Transportation Systems, Proceedings, ITSC*, vol. 2005, 2005, pp. 907–912.

[8] C. Miyajima, Y. Nishiwaki, K. Ozawa, T. Wakita, K. Itou, and K. Takeda, "Cepstral Analysis of Driving Behavioral Signals for Driver Identification," *2006 IEEE International Conference on Acoustics Speech and Signal Processing Proceedings*, vol. 5, no. October 2015, pp. 6–9, 2006.

[9] C. Miyajima, Y. Nishiwaki, K. Ozawa, T. Wakita, K. Itou, K. Takeda, and F. Itakura, "Driver modeling based on driving behavior and its evaluation in driver identification," *Proceedings of the IEEE*, vol. 95, no. 2, pp. 427–437, 2007.

[10] M. V. Martínez, J. Echanobe, and I. del Campo, "Driver identification and impostor detection based on driving behavior signals," pp. 372–378, 2016.

[11] A. Riener and A. Ferscha, "Supporting Implicit Human-to-Vehicle Interaction : Driver Identification from Sitting Postures," in *Proceedings of the First Annual International Symposium on Vehicular Computing Systems. July 22-24, 2008, Dublin*, 2008, pp. 1–10.

[12] J.-D. Wu and S. H. Ye, "Driver identification using finger-vein patterns with Radon transform and neural network," *Expert Systems with Applications*, vol. 36, no. 3 PART 2, pp. 5793–5799, 2009.

[13] J.-D. Wu and S.-H. Ye, "Driver identification based on voice signal using continuous wavelet transform and artificial neural network techniques," *Expert Systems with Applications*, vol. 36, no. 2, pp. 1061–1069, 2009.

[14] D. Hallac, A. Sharang, R. Stahlmann, A. Lamprecht, M. Huber, M. Roehder, R. Sosič, and J. Leskovec, "Driver identification using automobile sensor data from a single turn," pp. 953–958, 2016.

[15] S. Ezzini, I. Berrada, and M. Ghogho, "Who is behind the wheel? driver identification and fingerprinting," *Journal of Big Data*, vol. 5, no. 1, p. 9, Feb 2018. [Online]. Available: https://doi.org/10.1186/s40537-018-0118-7

[16] J. Principe, D. Xu, and J. Fisher, "Information-Theoretic Learning," *Unsupervised Adaptive Filtering, Volume 1: Blind Source Separation*, pp. 1–62, 2000.

[17] L. Breiman, "Random Forests," *Machine Learning*, vol. 45, no. 1, pp. 5–32, 2001.

[18] "Randomforest-matlab," https://github.com/jrderuiter/randomforest-matlab, 2013.

[19] S.-H. Park and J. Fürnkranz, "Efficient Pairwise Classification," in *Proceedings of the 18th European Conference on Machine Learning (ECML 2007, Warsaw, Poland)*, J. N. Kok, J. Koronacki, R. de Mántaras, S. Matwin, D. Mladenić, and A. Skowron, Eds. Springer-Verlag, 2007, pp. 658–665.