

Profiling Electric Vehicles via Early Charging Voltage Patterns

Francesco Marchiori¹[0000-0001-5282-0965], Denis Donadel²[0000-0002-7050-9369],
Alessandro Brighente¹[0000-0001-6138-2995], and
Mauro Conti¹[0000-0002-3612-1934]

¹ University of Padova, Padova, Italy

² University of Verona, Verona, Italy

francesco.marchiori@math.unipd.it, denis.donadel@univr.it,
{alessandro.brighente, mauro.conti}@unipd.it

Abstract. Electric Vehicles (EVs) are rapidly gaining adoption as a sustainable alternative to fuel-powered vehicles, making secure charging infrastructure essential. Despite traditional authentication protocols, recent results showed that attackers may steal energy through tailored relay attacks. One countermeasure is leveraging the EV's fingerprint on the current exchanged during charging. However, existing methods focus on the final charging stage, allowing malicious actors to consume substantial energy before being detected and repudiated. This underscores the need for earlier and more effective authentication methods to prevent unauthorized charging. Meanwhile, profiling raises privacy concerns, as uniquely identifying EVs through charging patterns could enable user tracking.

In this paper, we propose a framework for uniquely identifying EVs using physical measurements from the early charging stages. We hypothesize that voltage behavior early in the process exhibits similar characteristics to current behavior in later stages. By extracting features from early voltage measurements, we demonstrate the feasibility of EV profiling. Our approach improves existing methods by enabling faster and more reliable vehicle identification. We test our solution on a dataset of 7408 usable charges from 49 EVs, achieving up to 0.86 accuracy. Feature importance analysis shows that near-optimal performance is possible with just 10 key features, improving efficiency alongside our lightweight models. This research lays the foundation for a novel authentication factor while exposing potential privacy risks from unauthorized access to charging data.

Keywords: Electric Vehicles · Profiling · Voltage

1 Introduction

The increasing diffusion of Electric Vehicles (EVs), with total sales forecasted to reach up to 31.1 million by 2030 [10], is a key factor to help fight global warming. At the same time, it opens up to new challenges related to the peculiarity of these devices. Many of these challenges are related to the batteries, how to make a full

charge last longer, and how to charge them fast and efficiently, while ensuring the grid stability. The grid leverages communication between vehicles and charging columns by employing the so-called Vehicle-To-Grid (V2G) paradigm. Data exchange usually happens using a power line communication through the control pilot pin of the charging plug and is usually enabled by the ISO 15118 protocol [25]. Through it, EVs can establish a full-featured internet connection with the smart grid, allowing price negotiation, charging time scheduling, authentication handling, and enabling many other services [25]. While this communication comes with high benefits, it could also be exploited by attackers [2, 4, 9].

Being a Cyber-Physical System (CPS), the EV charging infrastructure security must not only rely on the protocol’s security but also consider threats from the physical world. In fact, information leakages can happen through the exchange of physical signals. In particular, researchers demonstrated how it is possible to profile an EV by looking at the energy delivered by the charging column [6, 7]. This approach collects the charging current in the final part of a full charging process to extract features and characterize each EV. The approach achieves good classification performance and is open to several potential implications, both for attackers who want to trace a vehicle between multiple locations and for defenders as an authentication mechanism. However, since the data are collected during the last part of the charging process (the so-called *tail*), the applicability of this approach is limited to charging processes not reaching 100% of State of Charge (SoC), thus potentially missing several profiling chances and making authentication more complicated.

Contributions. In this paper, we analyze a different solution that allows EV profiling by collecting measurements during the early stages of the charging process. In particular, we employ voltage measurements, which are significant in the first part of the charging process, instead of current values, which provide more information in the last part of the charging. Moreover, while similar works consider a dataset coming from a series of charging columns employing an adaptive charging solution (ACN Dataset [17, 18]), in this paper, we consider the EVBattery dataset [14, 33], which provides a more standard and spread charging management system. Overall, our findings demonstrate that EV authentication can be reliably performed using simple and lightweight models. Our approach achieves an accuracy of up to 0.86, averaged across all vehicle types.

The contribution of this paper can be summarized as follows:

- We analyze the feasibility of profiling a vehicle through measurements of the voltage exchange during the first charging steps.
- We analyze the performance of our framework on a large real-world dataset [33], comprising charges from more than 49 EVs and including three different brands. Our approach is tested on anonymized data, demonstrating its robustness in privacy-preserving scenarios.
- We identify the most relevant features for EV authentication, showing that a small subset of 10 key features is sufficient for high accuracy.
- We make our implementation and code open-source at: <https://github.com/spritz-group/EV-Volt-Auth>.

Organization. This paper is organized as follows. In Section 2 we report related works in the field of EV security and authentication. We then detail the considered system and threat model in Section 3, and in Section 4 we propose our methodology. We report the results of our evaluation in Section 5, and Section 6 concludes our work.

2 Related Works

Several research efforts have explored the security and privacy aspects of EVs [2, 13, 24]. Looking at the bigger picture, the need for EV to periodically connect to the grid for charging can threaten the entire power grid. Its stability could be mined by botnets of vehicles [16] and by coordinated attacks [12]. Moreover, the backend connection between control centers and the charging columns opens up several challenges [26]. For instance, the most widely used protocol for backend communications has been proved vulnerable to cyberattacks [1].

Being a CPS, the EV itself exhibits different security risks in addition to traditional petrol-powered cars [5, 31]. Attackers may exploit the physical connection to the charging columns to steal energy from a nearby victim [9] or perform denial of charge attacks, even remotely [4]. Different authentication strategies have been proposed to secure the charging process, investigating also the dynamic charging that allows cars to charge while moving [3, 19, 23]. Brighente et al. [7] first introduced an approach for authentication by employing information from the EV specific charging pattern, which was then expanded in different following works [6, 11]. However, their approach requires a full charge of the EV battery before producing a result, thus reducing the possible applicability of the framework. The same issue holds for authentication methods aimed at lithium-ion batteries, as they can require full battery charges and discharges or sophisticated equipment, making their usability in the context of EVs limited [21].

Privacy aspects of EV charging have also been investigated [30]. However, it is not always easy to prevent CPS from leaking information that an attacker may exploit to mine the user’s privacy [32]. Recent work has demonstrated the feasibility of extracting sensitive information—such as user identity, driving style, and trip endpoints purely from battery consumption patterns using Machine Learning (ML) techniques [22]. While this highlights the privacy risks inherent in battery-related telemetry, such approaches focus on post-drive consumption data, typically requiring access to a full trip. In contrast, our work shifts the focus to the early stages of the charging process, showing that brief voltage readings alone can be leveraged for profiling, even without full charging sessions or vehicle usage data. This significantly broadens the threat landscape, as it reduces the time and access requirements for potential adversaries.

3 System and Threat Model

In this section, we describe the system under study (Section 3.1), outline the adversarial model (Section 3.2), and explore potential use cases of our profiling

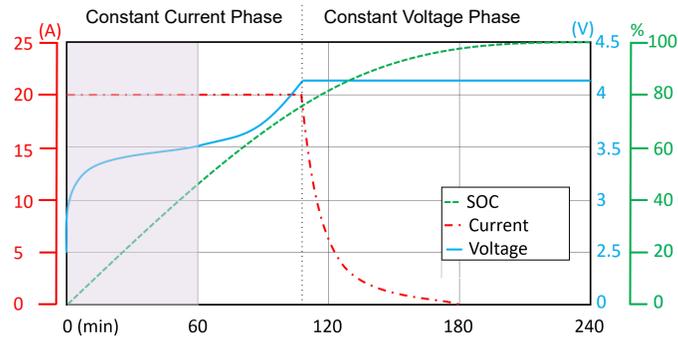


Fig. 1: Charging profile of a Li-ion battery [28]. In this paper, we will show that voltage data from the first part of the constant current phase (highlighted in violet in the graph) are enough for correctly profiling a vehicle.

method (Section 3.3). We distinguish between malicious and legitimate applications, highlighting the dual-use nature of EV fingerprinting via voltage traces.

3.1 System Model

Charging an EV requires careful consideration of both the vehicle and the grid. This latter, in particular, must be managed to handle high power requirements from a fleet of EVs, and this is facilitated by the V2G paradigm that creates a bi-directional communication between vehicles and the smart grid. On the other side, EV's battery can be charged at different speeds based on the capabilities of the charging column and the vehicle itself [15].

EVs are usually equipped with batteries containing Lithium-ion (Li-ion) cells. They exhibit a known and peculiar charging pattern where voltage and current levels are correlated to the SoC. In particular, two main behaviors can be identified, as shown in Fig. 1. The charging starts with a *constant current phase* where the voltage slowly increases. At a certain point, usually with SoC values between 60% and 80% of the full charge, a *constant voltage phase* starts, where the current slowly decreases, creating a tail. This behavior is well known and is typical of these kinds of batteries [27].

3.2 Threat Model

In this paper, we consider an entity that can collect measurements during the charging process. In particular, with respect to previous works [6] that collected current levels, we collect voltage levels during the charging process. As we discuss in this paper, this reduces the required collection time to profile a vehicle and does not require a full charge to extract the relevant features. The only other information required by the system is the SoC of the battery, which is used to understand in which part of the charging profile the EV is. All the other data

transmitted through the cable, such as the high-level communication transmitted in the control pilot pin, is not collected for this study.

Such profiling could be a double-edged sword and can be applied by different entities for different purposes. An attacker may exploit this methodology to track a vehicle through different locations. A malicious owner of various parking lots with charging columns may track vehicles for advertisement purposes. Another option for an attacker is to install a measurement device as a plug, similar to devices employed in ATM skimming [29], to collect data useful to track a user mining their privacy. On the other side, such data could be employed for good as a second factor or continuous authentication system.

3.3 Applications

We show that it is possible to extract discriminative features from early charging voltage patterns, which are unique enough to allow reliable profiling of a singular vehicle. The ability to fingerprint an EV from voltage measurements can be leveraged in multiple contexts. We highlight both adversarial and defensive uses.

Adversarial Applications.

- **Vehicle tracking across locations:** A malicious entity operating multiple public or semi-public charging stations can silently collect voltage traces and match them across sites to track vehicles over time.
- **Profiling without consent:** An attacker could embed a skimming device within a charging cable or socket to passively collect voltage traces and associate them with a specific vehicle or user, compromising location privacy.
- **Behavioral surveillance:** By linking profiling data with usage patterns, an attacker might infer sensitive behavior (e.g., commuting habits, work location, or home address).

Defensive Applications.

- **Second-factor authentication:** EV profiling can supplement existing user authentication mechanisms, confirming vehicle identity as part of a multi-factor security scheme at high-security charging stations.
- **Tamper detection:** Deviations in the expected voltage pattern could be used to detect unauthorized battery replacement or tampering with internal battery components.
- **Anti-theft tracking:** In case of theft, a known voltage fingerprint can serve as a unique signature to detect the vehicle if it is connected to any charging infrastructure.
- **Usage-based insurance or leasing:** Insurance or leasing companies may use this approach to verify the EV’s identity under a specific contract, enforcing user-vehicle binding in shared or rental contexts.

These use cases demonstrate that early-stage voltage profiling has a wide range of implications, from privacy threats to enabling lightweight security mechanisms. Our work does not advocate for any particular application but aims to

provide a technical foundation and empirical validation for EV identification based on early charging behavior. Depending on who controls the charging infrastructure, this capability may pose a risk to user privacy or be a valuable tool for improving EV security and authentication.

4 Methodology

In this section, we provide a more detailed explanation of the profiling techniques we employ. We first present the dataset and its characteristics (Section 4.1). We also discuss our feature extraction process, which is one of the key components of our approach (Section 4.2). Next, we present an overview of the ML models used as classifiers and the process of finding their optimal hyperparameters (Section 4.3). The overall methodology is summarized in Fig. 2. The first step is collecting charging measurements from EV charging process. From the samples, feature extraction is used to obtain information that can be fed to ML models for classification.



Fig. 2: Pipeline of the proposed profiling framework.

4.1 Dataset

For this work, we created a dataset starting from three EV charging collections, which together consist of over 690,000 charging snippets recorded from 347 distinct EVs [33]. Since the original dataset was designed for anomaly detection, we first remove all potential outliers that exhibit anomalies in the charging process. Next, we extract all charging snippets, which originally are segmented into 128-sample sequences, with each sample recorded once per second. Our primary feature of interest is voltage, though we also incorporate SoC for further processing. To ensure we analyze complete charging events, we concatenate snippets belonging to the same charging session, grouping them based on the car label and segment ID. This process results in complete charge sequences, which we sort by SoC to maintain its natural increasing order. After preprocessing, our dataset consists of voltage and SoC time series, and car labels. From this, we apply two additional filtering steps.

- *SoC Thresholding* – We keep only data where SoC is $\leq 60\%$, as our profiling approach relies on the non-constant voltage phase in the early charging stages (see Fig. 1).

- *Minimum Sample Requirement* – We include only vehicles with at least 100 charging samples, ensuring sufficient data for profiling. Cars with few charging instances are excluded due to insufficient label representation in the dataset.

The final dataset thus contains 36,165 charging snippets constituting 7,408 charging sessions from 49 different EVs. It is worth noting that the voltage data has been perturbed and interpolated as part of the anonymization process applied by the original dataset authors [33]. Despite these modifications, underlying temporal patterns and correlations remain intact, potentially enabling the inference of an EV’s identity, an aspect we explore further in Section 5.5. After this one-time preprocessing step, we further adapt the dataset dynamically to be used in a classification setup. The final distribution of the dataset is shown in Fig. 3. From now on, we will discuss an authentication scenario, but a malicious user could apply the same process to profile and track a vehicle between charging stations. The only subtle difference is that an attacker may not have access to detailed SoC data since the information could be transmitted encrypted or on channels not under the attacker’s control. We will discuss this issue in Section 5.3.

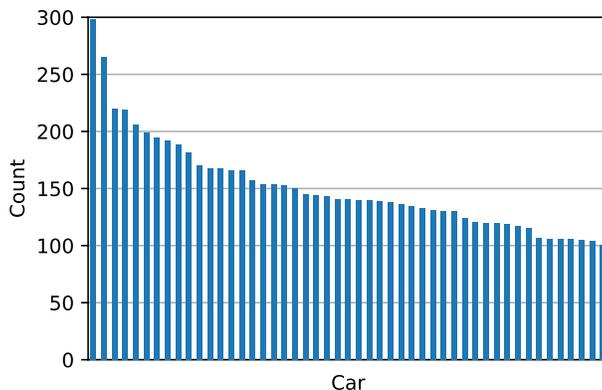


Fig. 3: Distribution of samples for each car in our dataset after pre-processing.

In particular, we set up our systems to first select a vehicle as the authenticated subject for each car label and consider all other vehicles as non-authenticated. This process is repeated for every vehicle in the dataset, ensuring that each car is evaluated as the authenticated one at least once. By structuring the dataset this way, we create multiple binary classification tasks tailored to a specific vehicle’s charging behavior.

4.2 Feature Extraction

For each time window, we extract features from the voltage and SoC time series using the `tsfresh` Python package [8]. `tsfresh` automatically computes a wide range of statistical and mathematical characteristics from time series data and selects the most relevant ones for regression or classification tasks. The extracted features include but are not limited to, statistical measures (e.g., mean, variance, skewness, kurtosis), frequency-domain features (e.g., Fourier coefficients, spectral entropy), and time-series properties (e.g., autocorrelation, trend strength, peaks, crossings). We apply this process to the voltage and SoC time series from our processed dataset, ensuring that each driver performs feature extraction separately. This allows us to generate feature representations tailored to each driver’s charging behavior, which we then use in a binary classification setting (authenticated vs. non-authenticated). After extracting features, we perform two processing steps.

1. *Feature Imputation* – Some extracted features may contain missing values due to insufficient data in certain time windows. We handle this by imputing missing values, typically by replacing them with appropriate statistical estimates (e.g., mean, median, or interpolation methods), ensuring a complete dataset for classification.
2. *Feature Selection* – Since `tsfresh` generates many features, we apply feature selection to retain only the most informative and discriminative ones. The reduction of feature number is done independently for each EV to maintain the most suited features in each case. This step reduces dimensionality and improves model efficiency, removing noisy or redundant features. Although feature selection is performed independently for each vehicle, we observe a significant overlap in the selected features across different models. This suggests that certain statistical patterns are consistently informative, regardless of the individual EV’s characteristics. We further investigate this in Section 5.4, where we analyze feature importance and demonstrate that a small set of just 10 common features is sufficient to achieve performance close to that of the full feature set.

4.3 Models

Inspired by previous work on EV authentication [6] and battery authentication [21], we select five lightweight ML models: AdaBoost, Decision Tree (DT), k-Nearest Neighbors (kNN), Neural Network (NN), and Random Forest (RF). These models were chosen not only based on their effectiveness in prior studies but also due to their suitability for real-time deployment. Unlike deep learning approaches, which typically require high computational power and external data processing, these models are lightweight enough to run directly on an EV charging station without requiring communication with external servers, enhancing security by reducing potential attack vectors associated with remote authentication. We apply an 80/20 split between the training and test sets for model

training. Hyperparameter tuning is performed using a grid search approach with 5-fold cross-validation, ensuring that our models generalize well to unseen data. The set of hyperparameters considered for each model is detailed in the Appendix A.1.

5 Evaluation

We now proceed with the evaluation of our models under different settings. In particular, we analyzed the behavior of our models when changing the measurement length to extract the features from (Section 5.1) and when varying the composition of the dataset (Section 5.2). Moreover, we analyzed the importance of the SoC feature in Section 5.3, the features employed and the effects on reducing the feature number in Section 5.4.

To evaluate our models, we employ two classical metrics, the accuracy and the F1-score, which are defined as follows:

$$Accuracy = \frac{TP + TN}{TP + FP + TN + FN}, \quad F1 = \frac{2TP}{2TP + FP + FN}, \quad (1)$$

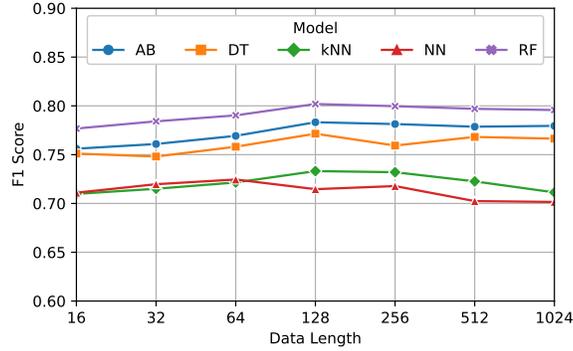
where:

- *True Positive (TP)*: a sample from the authenticated vehicle is correctly identified.
- *False Positive (FP)*: a sample belonging to the non-authenticated class mistakenly identified as the authenticated vehicle.
- *False Negative (FN)*: a sample belonging to the authenticated driver is mistakenly identified as a non-authenticated driver.
- *True Negative (TN)*: a non-authenticated vehicle is correctly classified as non-authenticated.

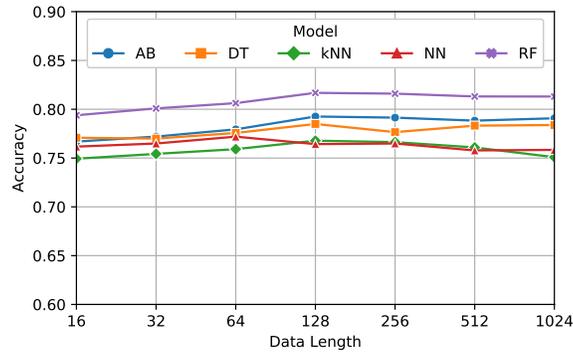
5.1 Data Length

One of the most important differences with respect to previous works [6] is that our models do not require a full charge to efficiently profile a vehicle. In fact, since we extract features from the first part of the charging, only a few seconds of charging are sufficient to obtain some results. To understand how much the size of the data capture has an impact on the performance, we tested several options ranging from 16 to 1024 seconds.

Results are shown in Fig. 4, where Fig. 4a shows the F1-scores, while Fig. 4b the accuracy, both averaged by the ratios (see Section 5.2). As we can see, the overall variance between sizes is not very pronounced, and this suggests it is not necessary to wait long to make a decision. In particular, top scores are reached with a 128-second length by the RF classifier, which overall performed quite better than other models. Another interesting aspect is related to the results of different models within the same data length. The RF models retain an almost



(a) F1 Score of the models.



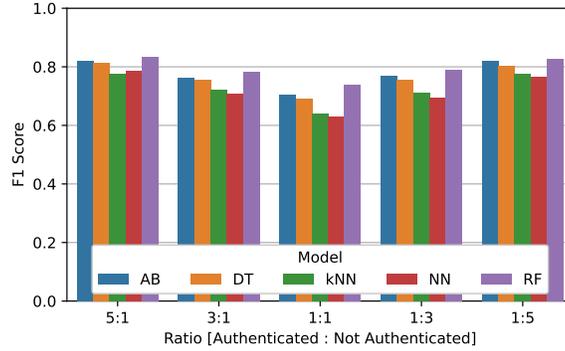
(b) Accuracy of the models.

Fig. 4: Average performance of different data lengths (in seconds).

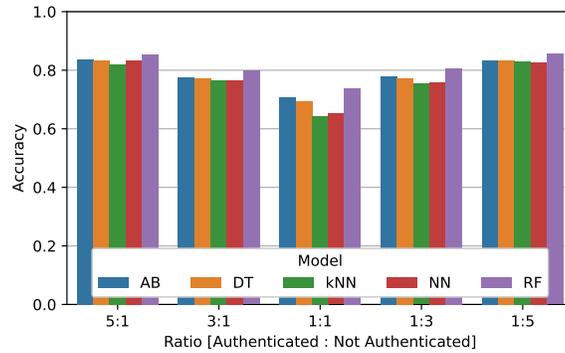
constant distance from the second most performing model over all the data lengths. Although kNN and NN represent two very different architectures, they share a curve that sees rising scores up to mid-length (i.e., 64 to 128) and then falling again as length increases.

5.2 Ratio

We conduct several experiments to verify the effect of an unbalanced dataset during training. To analyze it, we tested our models based on different ratios between the data concerning the positive label versus all the others. In particular, a ratio of 5 : 1 indicates that the samples belonging to the selected EV are five times more than samples from all the other vehicles. Conversely, a ratio of 1 : 5 indicates that the dataset contains five times more samples of other vehicles with respect to the selected one. A ratio of 1 : 1 indicates a balanced dataset.



(a) F1 Score of the models.



(b) Accuracy of the models.

Fig. 5: Average performance of different ratios (representing the number of authenticated samples vs the number of other samples in the dataset).

The results shown in Fig. 5 clearly indicate unbalanced datasets’ capabilities to perform successfully. This is essential in such a scenario where it is usually more difficult to get samples from the authenticated vehicle with respect to data from all the other cars that use the charging column. In particular, we can see how with one-fifth of data related to the target (i.e., ratio 1 : 5), the F1-score surpasses 0.80 for three models out of five. Even in this experiment, RF exhibits the best performances. However, the gap is more pronounced with a balanced dataset, while AdaBoost (AB) almost reaches the F1-Score levels of RF for heavy unbalanced datasets.

5.3 SoC Feature

In our experiments we employ the SoC as a time series for feature extraction through `tsfresh`. This is reasonable since we discussed an authentication sce-

nario where this information should be available to the charging column with a high level of detail. However, in an adversarial scenario, an attacker may only be able to get rough information about the SoC for instance by looking at the charging column display, or by estimating it from the voltage level and the duration of the charging. In this scenario, the attacker could successfully identify if the charging is happening during the constant current phase (see Fig. 1), but cannot use the SoC to extract features. To determine the impact of the absence of the SoC feature, we reproduced our experiments without considering it during the feature extraction, but employing it only to filter out samples with more than 60% of SoC. The results of this analysis are shown in Fig. 6.

As we can see, results are comparable, although the presence of SoC-related features slightly improve the performances of our classifier. Interestingly, we notice a more significant gap in accuracy performance as the considered data length for classification decreases. This suggests that SoC-related features help the model more accurately assess the specific charging phase. However, the model can infer this independently with more considerable data lengths, as suggested by the smaller gap between the accuracy results.

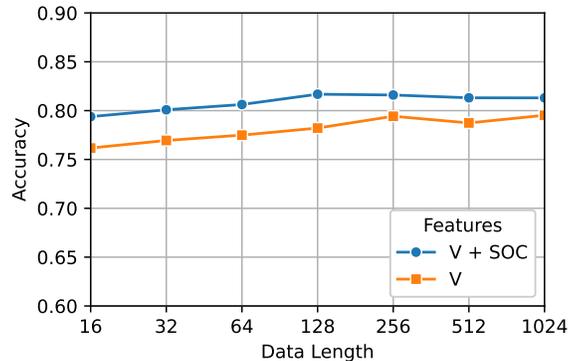


Fig. 6: Average performance of our RF model trained on different subsets of features (**V**oltage alone or with **SoC**) at different data lengths.

5.4 Feature Analysis

A complete analysis of the features extracted and employed by `tsfresh` is useful to understand which trait of the charging is more characteristic and thus useful in the profiling. We performed our feature analysis on a subsample of the whole datasets and employed the best-performing model, i.e., the RF model.

We start by analyzing the feature’s importance by employing SHapley Additive exPlanations (SHAP), a model-agnostic XAI technique [20]. Fig. 7 classify

the main features with their importance level, which are explained more into detail in Appendix A.2. The prominence of frequency-domain features among the top 10 highlights the significance of periodic voltage signal behavior in EV authentication. The inclusion of both low- and high-frequency components suggests that variations in power electronics and charging circuit dynamics are distinctive to individual EVs. Additionally, the presence of features such as mean absolute change and standard deviation indicates that, while absolute voltage levels may be comparable across EVs, their dynamic fluctuations provide a unique signature for identification.

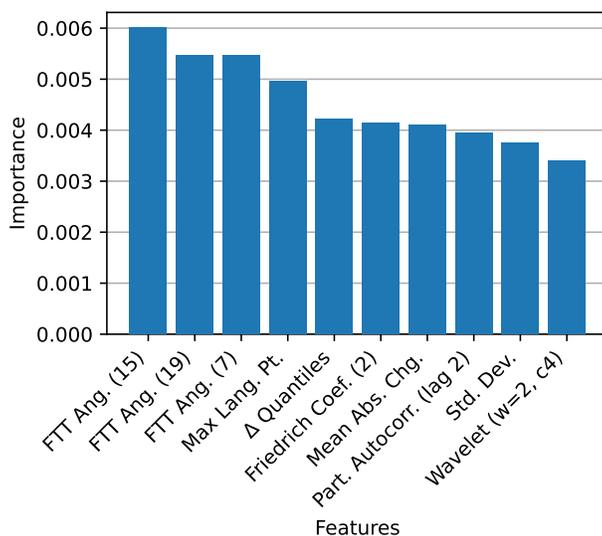


Fig. 7: Top 10 features by importance extracted by SHAP.

After understanding the most significant features, another important aspect is the impact of adding features on the scores. A low number of features reduces the complexity of the model, making it more easy to handle and less exposed to noise. Therefore, we experiment with different numbers of features in a RF model, choosing every time the x most significant features. Results are shown in Fig. 8. They show that increasing the number of features generally improves accuracy, but the rate of improvement diminishes as more features are added, becoming almost negligible from 6 features on. Moreover, the effect varies depending on the data balance ratio. Higher imbalance ratios tend to achieve better performance overall, suggesting that the model benefits from more training samples in the majority class. Even with balanced datasets, accuracy improves until plateauing, showing that a feature subset can achieve near-optimal results while keeping the model efficient.

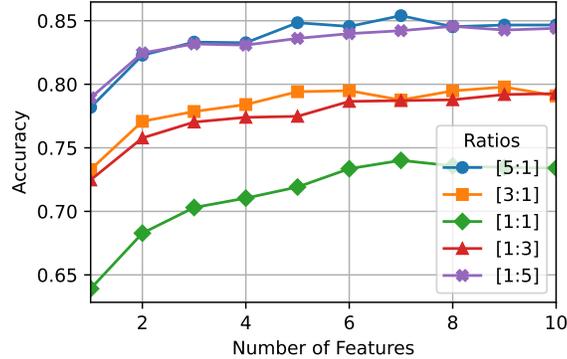


Fig. 8: Accuracy of the classifiers trained on different dataset ratios (i.e., number of authenticated samples vs number of other samples) while increasing the number of features.

5.5 Data Anonymization

The authors of the original dataset [33] applied a series of anonymization steps to protect the identity and characteristics of the EVs. Specifically, they perturbed and interpolated each charging segment’s average voltage, current, and temperature values. Additionally, timestamps and mileage values were randomly shifted and scaled. While these anonymization techniques help mask exact values and protect sensitive identifiers, our results show that meaningful patterns remain. In particular, the temporal correlations and statistical structures in the voltage data can effectively be leveraged to infer the EV’s identity. This demonstrates the robustness of our methodology not only in idealized settings but also when applied to anonymized real-world data. This highlights the effectiveness of our methodology in both system and threat model scenarios.

- *Benign use case*: When employed for authentication at the charging station level, our method enables privacy-preserving user recognition resilient against side-channel attacks on battery usage [22].
- *Adversarial use case*: When used for malicious profiling, our findings indicate that the current anonymization strategy may be insufficient, pointing to stronger or alternative techniques to prevent identity inference.

6 Conclusions

In this paper, we analyzed the feasibility of profiling an EV during charging by exploiting voltage measurements and the SoC. The data employed are collected during the first part of the charging process, thus allowing vehicle identification during the first stages of the charging, without the need to wait for full charging

(i.e., 100% SoC). Our experiments on 49 different vehicles show accuracy up to 0.86, demonstrating the practicality of the approach.

Privacy is an essential requirement in the CPS environment. Looking at potential leakages in the EV scenario is important to defend users against possible misuse of such information. Moreover, in this context, this information could also be employed to enhance the security of the system, providing novel features for authentication. In fact, we foresee future work that can build a fully featured second-factor and continuous authentication system starting from this paper’s findings. Using advanced preprocessing can also enhance scores and enable more accurate, robust profiling. Moreover, a larger dataset could allow for investigating the resiliency of the solution to physical properties such as battery aging, seasonality, or time of day.

Open Science and Ethical Considerations

This work builds upon publicly available, pre-processed EV charging data released for research purposes. While raw data remain protected under privacy regulations, the processed dataset includes vehicle identifiers that allow charging sessions to be linked to specific (anonymized) EVs. Our findings demonstrate that even with such limited information, accurate profiling remains feasible—highlighting a broader security and privacy consideration for the community. We believe it is important to share these results to help guide the development of more privacy-preserving data-sharing practices and protocols in future EV infrastructures. All our code is available to promote transparency and reproducibility: <https://github.com/spritz-group/EV-Volt-Auth>.

Acknowledgment

This work was funded by the European Union under the National Recovery and Resilience Plan (NRRP), Mission 4 Component 2 Investment 1.3 - Call for tender No. 341 of March 15, 2022 of Italian Ministry of University and Research – NextGenerationEU; Code PE00000014, Concession Decree No. 1556 of October 11, 2022 CUP D43C22003050001, Project “SEcurity and RIghts in the CyberSpace (SERICS) - Spoke 7 Infrastructure Security - Visible Light Communication for Secure Vehicle-to-Everything Communication - VisiCar” – Beneficiary’s CUP: C99J24000250008.

References

1. Alcaraz, C., Cumplido, J., Trivino, A.: Ocpc in the spotlight: threats and countermeasures for electric vehicle charging infrastructures 4.0. *International Journal of Information Security* **22**(5), 1395–1421 (2023)
2. Antoun, J., Kabir, M.E., Moussa, B., Atallah, R., Assi, C.: A detailed security assessment of the ev charging ecosystem. *IEEE Network* **34**(3), 200–207 (2020)

3. Babu, P.R., Amin, R., Reddy, A.G., Das, A.K., Susilo, W., Park, Y.: Robust authentication protocol for dynamic charging system of electric vehicles. *IEEE Transactions on Vehicular Technology* **70**(11), 11338–11351 (2021)
4. Baker, R., Martinovic, I.: Losing the car keys: Wireless PHY-Layer insecurity in EV charging. In: 28th USENIX Security Symposium (USENIX Security 19). pp. 407–424. USENIX Association, Santa Clara, CA (Aug 2019), <https://www.usenix.org/conference/usenixsecurity19/presentation/baker>
5. Brighente, A., Conti, M., Donadel, D., Poovendran, R., Turrin, F., Zhou, J.: Electric vehicles security and privacy: Challenges, solutions, and future needs. arXiv preprint arXiv:2301.04587 (2023)
6. Brighente, A., Conti, M., Donadel, D., Turrin, F.: Evscout2. 0: Electric vehicle profiling through charging profile. *ACM Transactions on Cyber-Physical Systems* **8**(2), 1–24 (2024)
7. Brighente, A., Conti, M., Sadaf, I.: Tell me how you re-charge, i will tell you where you drove to: Electric vehicles profiling based on charging-current demand. In: Bertino, E., Shulman, H., Waidner, M. (eds.) *Computer Security – ESORICS 2021*. pp. 651–667. Springer International Publishing, Cham (2021)
8. Christ, M., Braun, N., Neuffer, J., Kempa-Liehr, A.W.: Time Series Feature Extraction on basis of Scalable Hypothesis tests (tsfresh – A Python package). *Neurocomputing* **307**, 72–77 (2018). <https://doi.org/https://doi.org/10.1016/j.neucom.2018.03.067>
9. Conti, M., Donadel, D., Poovendran, R., Turrin, F.: Evexchange: A relay attack on electric vehicle charging system. In: *European Symposium on Research in Computer Security*. pp. 488–508. Springer (2022)
10. Deloitte: Electric vehicles: setting a course for 2030. In: <https://www2.deloitte.com/us/en/insights/focus/future-of-mobility/electric-vehicle-trends-2030.html> (Jul 2020)
11. Gangwal, A., Jain, A., Conti, M., et al.: On the feasibility of profiling electric vehicles through charging data. In: *Inaugural International Symposium on Vehicle Security & Privacy* (2023)
12. Ghafouri, M., Kabir, E., Moussa, B., Assi, C.: Coordinated charging and discharging of electric vehicles: A new class of switching attacks. *ACM Transactions on Cyber-Physical Systems (TCPS)* **6**(3), 1–26 (2022)
13. Gottumukkala, R., Merchant, R., Tauzin, A., Leon, K., Roche, A., Darby, P.: Cyber-physical system security of vehicle charging stations. In: *2019 IEEE Green Technologies Conference(GreenTech)*. pp. 1–5 (2019)
14. He, H., Zhang, J., Wang, Y., Jiang, B., Huang, S., Wang, C., Zhang, Y., Xiong, G., Han, X., Guo, D., et al.: Evbattery: A large-scale electric vehicle dataset for battery health and capacity estimation. arXiv preprint arXiv:2201.12358 (2022)
15. Khalid, M.R., Khan, I.A., Hameed, S., Asghar, M.S.J., Ro, J.: A comprehensive review on structural topologies, power levels, energy storage systems, and standards for electric vehicle charging stations and their impacts on grid. *IEEE access* **9**, 128069–128094 (2021)
16. Khan, O.G.M., El-Saadany, E., Youssef, A., Shaaban, M.: Impact of electric vehicles botnets on the power grid. In: *2019 IEEE Electrical Power and Energy Conference (EPEC)*. pp. 1–5. IEEE (2019)
17. Lee, Z.J., Johansson, D., Low, S.H.: ACN-Sim: An open-source simulator for data-driven electric vehicle charging research. *2019 IEEE International Conference on Communications, Control, and Computing Technologies for Smart Grids, Smart-GridComm 2019* pp. 411–412 (2019). <https://doi.org/10.1109/SmartGridComm.2019.8909765>

18. Lee, Z.J., Li, T., Low, S.H.: ACN-Data: Analysis and Applications of an Open EV Charging Dataset. In: Proceedings of the Tenth ACM International Conference on Future Energy Systems. p. 139–149. e-Energy '19, Association for Computing Machinery, New York, NY, USA (2019)
19. Li, H., Dán, G., Nahrstedt, K.: Portunes+: Privacy-preserving fast authentication for dynamic electric vehicle charging. *IEEE Transactions on Smart Grid* **8**(5), 2305–2313 (2016)
20. Lundberg, S.M., Lee, S.I.: A unified approach to interpreting model predictions. In: Guyon, I., Luxburg, U.V., Bengio, S., Wallach, H., Fergus, R., Vishwanathan, S., Garnett, R. (eds.) *Advances in Neural Information Processing Systems* 30, pp. 4765–4774. Curran Associates, Inc. (2017), <http://papers.nips.cc/paper/7062-a-unified-approach-to-interpreting-model-predictions.pdf>
21. Marchiori, F., Conti, M.: Your battery is a blast! safeguarding against counterfeit batteries with authentication. In: Proceedings of the 2023 ACM SIGSAC Conference on Computer and Communications Security. pp. 105–119 (2023)
22. Marchiori, F., Conti, M.: Leaky batteries: A novel set of side-channel attacks on electric vehicles. arXiv preprint arXiv:2503.08956 (2025)
23. Mookherji, S., Odelu, V., Prasath, R.: Secure ultra fast authentication protocol for electric vehicle charging. *Computers and Electrical Engineering* **119**, 109512 (2024)
24. Muhammad, Z., Anwar, Z., Saleem, B., Shahid, J.: Emerging cybersecurity and privacy threats to electric vehicles and their impact on human and environmental sustainability. *Energies* **16**(3), 1113 (2023)
25. Mültin, M.: Iso 15118 as the enabler of vehicle-to-grid applications. In: 2018 International Conference of Electrical and Electronic Technologies for Automotive. pp. 1–6 (2018)
26. Nasr, T., Torabi, S., Bou-Harb, E., Fachkha, C., Assi, C.: Chargeprint: A framework for internet-scale discovery and security analysis of ev charging management systems. In: NDSS (2023)
27. Shen, W., Vo, T.T., Kapoor, A.: Charging algorithms of lithium-ion batteries: An overview. In: 2012 7th IEEE conference on industrial electronics and applications (ICIEA). pp. 1567–1572. IEEE (2012)
28. ThunderSky: Instruction manual for LFP/LCP/LMP lithium power battery. Tech. rep., Thunder Sky (2007)
29. United States Secret Service: ATM & POS Skimming. <https://www.secretservice.gov/investigations/skimming> (2024)
30. Unterweger, A., Knirsch, F., Engel, D., Musikhina, D., Alyousef, A., de Meer, H.: An analysis of privacy preservation in electric vehicle charging. *Energy Informatics* **5**(1), 3 (2022)
31. Ye, J., Guo, L., Yang, B., Li, F., Du, L., Guan, L., Song, W.: Cyber-physical security of powertrain systems in modern electric vehicles: Vulnerabilities, challenges, and future visions. *IEEE Journal of Emerging and Selected Topics in Power Electronics* **9**(4), 4639–4657 (2020)
32. Zhang, H., Shu, Y., Cheng, P., Chen, J.: Privacy and performance trade-off in cyber-physical systems. *IEEE Network* **30**(2), 62–66 (2016)
33. Zhang, J., Wang, Y., Jiang, B., He, H., Huang, S., Wang, C., Zhang, Y., Han, X., Guo, D., He, G., et al.: Realistic fault detection of li-ion battery via dynamical deep learning. *Nature Communications* **14**(1), 5940 (2023)

A Appendix

A.1 Hyperparameters

Table 1 lists the models employed in classification and their hyperparameters.

Table 1: Hyperparameters employed in Grid Search.

Models	Hyperparameters
AdaBoost (AB)	<ul style="list-style-type: none"> • Number of estimators
Decision Tree (DT)	<ul style="list-style-type: none"> • Criterion • Maximum Depth
k-Nearest Neighbors (kNN)	<ul style="list-style-type: none"> • Number of neighbors • Weight function
Neural Network (NN)	<ul style="list-style-type: none"> • Hidden layer sizes • Activation function • Solver
Random Forest (RF)	<ul style="list-style-type: none"> • Criterion • Number of estimators

A.2 Feature Importance

The 10 most important features for the classification of the legitimate EV are the following:

1. *FFT angle (coeff. 15)* – The phase angle of the 15th Fourier coefficient, capturing periodic patterns in the voltage signal.
2. *FFT angle (coeff. 99)* – Similar to the previous feature but for the 99th Fourier coefficient, highlighting high-frequency behaviors.
3. *FFT angle (coeff. 7)* – The phase angle of the 7th Fourier coefficient, indicative of lower frequency components.
4. *Max Langevin point* – A stability indicator derived from stochastic differential equations, estimating system behavior.
5. *Change quantiles (mean)* – Measures the mean absolute change between specific quantiles, capturing shifts in voltage distribution.
6. *Friedrich coefficient (2)* – A coefficient from Friedrich’s method, modeling the dynamics of the voltage signal.
7. *Mean absolute change* – Computes the mean of absolute differences between consecutive voltage values, indicating signal variability.
8. *Partial autocorrelation (lag 2)* – Measures the correlation of the voltage signal with its past values at a lag of 2, detecting dependencies.
9. *Standard deviation* – Quantifies the overall dispersion and variation of the voltage signal.
10. *Wavelet coeff. (w=2, coeff. 4)* – Extracted from continuous wavelet transform, capturing multi-scale fluctuations in the voltage signal.