

Electric Vehicle Charging Cybersecurity from a Power Grid: A Review

Vedika Nerkar¹, Mrunmayi Pusrekar², Rahul Agrawal³, Sonali Gosavi⁴

¹²³⁴Department of Electrical Engineering, Guru Gobind Singh College of Engineering and Research Centre, Nashik, Maharashtra, India

Abstract - Electric vehicles, charging stations and the power grid now form a closely coupled cyber-physical system, in which a flaw in any one layer can be used to disrupt the others. Taking the grid-focused threat analysis of Acharya et al. as its foundation, this paper tests it against five recent studies recurrent-network anomaly detection on charger traffic, fuzzy-neural classification of in-vehicle bus traffic and three literature surveys together with five further peer-reviewed works on charger vulnerabilities, grid impact, management systems and mobile applications. A second study takes its class labels from the very identifier field that it also feeds in as an input feature, which may overstate accuracy. Across all sources the evidence relies chiefly on simulations, small test sets or single testbeds, and no work connects charger-level detection with grid-level consequences. The paper sketches a conceptual layered monitoring architecture and a prioritised list of research gaps. No new experiments are reported, and the proposed architecture has not been validated.

Key Words: Electric vehicle charging, Cybersecurity, Power grid, Anomaly detection, OCPP, CAN bus, Demand-side attack, STRIDE.

1. INTRODUCTION

Electric vehicles (EVs) have ceased to be passive electricity consumers. They now exchange data with charging stations, vendor clouds, smartphone applications and, more and more, with utilities. Acharya et al. [1] report that the worldwide EV industry expanded by roughly 60 percent a year from 2014 to 2019, that the count of charging stations (EVCSs) reached 7.3 million by the end of 2019, and that chargers of up to 350 kW can refill a vehicle in less than 15 minutes. Every one of these developments introduces additional communication links, and every link may serve as an entry point for an attacker. Since charging loads are large, concentrated and frequently controlled remotely, the consequences of a successful attack need not remain within the vehicle or the station.

This study builds on the paper by Acharya et al. [1]. That work models the vehicle, the charging station and the power grid as a single cyber-physical chain, lists the vulnerabilities found at each interface, and uses an attack tree to argue that compromised charging demand could be exploited to operate protection relays in the grid. Few studies bring these three perspectives together in one

threat model, which makes it a convenient reference point for assessing later research.

The literature has expanded rapidly since 2020. A systematic review of 102 publications describes a sharp increase in work on attacks and intrusion detection [2]; a wide-ranging survey screened 340 papers and kept 204 [3]; and another survey examines 102 AI-based intrusion detection proposals for in-vehicle networks [4]. Two studies from 2025 contribute concrete detection experiments, one on traffic between chargers and their controller [5] and one on controller area network (CAN) bus frames [6]. Because these works differ in data, metrics and threat assumptions, their results cannot be set side by side directly. The central question addressed here is therefore: to what extent do the published findings support the grid-level concern raised in [1], and what remains missing between device-level detection and grid-level consequences?

The contributions are as follows. (i) The cyber-physical model and threat model of the base paper are restated concisely and assessed. (ii) Five more peer-reviewed papers, whose bibliographic details were verified against publisher or indexing pages, are used to carry the base paper forward beyond 2020. (iii) A conceptual layered monitoring architecture and a set of research gaps are put forward. This paper is qualitative and reports no new experiments.

Section 2 describes the method and scope. Section 3 analyses the base paper. Section 4 appraises the five reference papers, and Section 5 presents the additional literature. Section 6 compares the evidence, Section 7 sets out the conceptual architecture and research gaps, Section 8 states the limitations, and Section 9 concludes.

2. METHOD AND SCOPE

The material comprises the base paper, five papers provided for critical analysis, and five additional papers located through publisher and indexing sources. These additional papers were selected using four criteria: peer-reviewed publication; a direct focus on EV or EVCS security or on grid impact; not being among the five provided papers; and coverage of attack surfaces that the provided papers address only briefly, namely vulnerability surveys, risk analysis, grid impact, charger management systems and mobile applications. For each of them, the title,

authors, venue and year were confirmed on a publisher, repository or indexing page.

For each paper, the analysis recorded the attack surface, assumptions, data, metrics and stated limitations. Attack paths were then assigned to the STRIDE categories used in [1]: spoofing, tampering, repudiation, information disclosure, denial of service and elevation of privilege. Wherever a paper gives a confusion matrix, the standard definitions of accuracy, precision, recall and F1 score were applied to it, and the outcome was compared with the figures given in the text. The comparison between studies is qualitative; it is not a systematic review, and no statistical pooling is attempted.

3. LITERATURE REVIEW

3.1 Cyber-physical structure

Acharya et al. [1] portray the EV as a collection of interconnected electronic control units (ECUs), of which modern vehicles hold more than 125. The ECUs exchange data over the CAN bus, the local interconnect network (LIN), media oriented systems transport (MOST) and FlexRay, all joined through a central gateway. The CAN bus predominates because it can handle simultaneous commands in real time, and for this reason it also carries the safety-critical functions. Its messages are broadcast to every node, are neither encrypted nor authenticated, and are ranked by an arbitration identifier. Limited bandwidth is cited as one reason why heavy cryptography has not been adopted. The on-board diagnostic (OBD-II) port offers direct access to the bus, and smartphone dongles can make it reachable wirelessly.

Outside the vehicle, the paper lists seven types of external interface: the charging station, physically accessible ports, internet service portals, radio and navigation signals, manufacturer update channels and roadside infrastructure. As regards charging, level 1 stations are wall outlets supplying 1.44 to 1.92 kW, level 2 stations provide single-phase 240 V or three-phase 400 V AC at up to 55 kW, and level 3 stations supply DC power of up to 350 kW. A DC station is portrayed as an AC station coupled with a high-power charger, which accounts for its larger software and control footprint. On the cyber side, a station includes an on-site human-machine interface (HMI), a connection to a building energy management system (BEMS) or demand-response aggregator via OpenADR, a connection to a central server, a smartphone interface and a vendor update channel. Fig -1 summarises this structure in four layers.




Fig -1: Layered representation of the vehicle–station–backend–grid chain adopted in this paper. The schematic is the author’s simplified version of the cyber-physical layers described in [1]; the arrows indicate the direction in which effects travel towards the grid.

3.2 Vulnerabilities and STRIDE classification

The base paper reviews documented weaknesses interface by interface and classifies the related literature using STRIDE. It also observes that smart charging infrastructure had not yet experienced a large, high-profile attack, although disclosed flaws in a smartphone application, a vendor charger line and a web application already demonstrated that remote access to chargers was realistic [1]. Table -1 summarises the principal findings; the STRIDE column is the author’s own mapping and not a table printed in the base paper.

Table -1 Interfaces, reported weaknesses and dominant STRIDE class (condensed from [1]; class mapping by the author)

Interface	Reported weakness	STRIDE class
In-vehicle CAN bus	Messages are not encrypted or authenticated and are delivered by broadcast; reachable through OBD-II	Spoofing, tampering, DoS
Tyre-pressure sensors	Static 32-bit identifiers sent unencrypted; spoofing possible from about 40 m away	Spoofing, disclosure
EV-station link (ISO 15118)	Charging messages intercepted at 54 stations using CCS connectors	Disclosure, spoofing
Station internals	Outdated Linux kernel, weak passwords, unsigned firmware; firmware can be extracted via debug and USB ports	Tampering, privilege
Station HMI	Exposed USB, serial and Ethernet ports; RFID credentials that can be cloned	Spoofing, tampering
Station-server link (OCPP)	Earlier versions lack certificates and end-to-end encryption; OCPP 2.0.1 introduced security profiles	Spoofing, tampering, repudiation
Apps and web portals	Credentials stored in plain text; authentication bypassed in disclosed cases	Spoofing, tampering
BEMS and demand response	Poor consumer-side practice; large IoT botnets as a precedent	DoS, tampering
Vendor update channels	Wireless updates can be intercepted and altered	Tampering, privilege

3.3 Grid threat model and reported impact

The main contribution of [1] is the attack tree, reproduced here in simplified form. Its root goal is an over- or under-frequency or voltage event that trips protection relays and may spread in a cascade. Two branches lead to this goal. In the first, a large number of stations that are in use are switched off simultaneously, causing demand to drop abruptly; the attacker either builds a botnet of stations (by infecting vehicles that then pass on malware, or by injecting malware directly) or compromises a station server. In the second, false availability, capacity or price data are broadcast to vehicles so that their demand is steered towards a weak part of the grid, with a station server, a fleet server or a smartphone application as the entry points. The paper emphasises that an attacker requires only publicly available station and grid data to decide which networks are worth compromising.

Table -2 lists the impact studies summarised in the base paper, along with one later study. The values are those given by the original authors, quoted through [1], and they originate from modelling or laboratory work rather than from field incidents.

Table -2 Reported grid-level effects of attacks on EV charging

Study	Attack studied	Reported effect
Acharya et al., as summarised in [1]	Manipulation of charging demand at public stations; Manhattan case constructed from public data	Load equal to about 1,000 EVs at 350 kW could activate over-frequency relays and cause major outages
Rohde, as summarised in [1]	Attack on the control system of a 50 kW DC station	Current distortion exceeding 20 percent and power factor under 0.8; the IEEE 519 limit quoted is 8 percent
Khan et al., as summarised in [1]	Botnet of EVs and fast-charging stations tested on IEEE 33- and 39-bus networks	Line congestion, under-voltage and outages; assumes an omniscient attacker
Morrison, as summarised in [1]	Botnet of 7 kW home chargers in California	About 12 percent of EVs switched off simultaneously, producing a 0.5 Hz frequency drop
Sayed et al. [7]	Load-altering attack using EV loads compared with residential loads	The test system was destabilised by a 30 MW EV attack but not by a 48 MW residential attack

3.4 Strengths of the base paper

The base paper's strengths are its integrated three-tier view, its use of a recognised threat taxonomy, and its feasibility argument built on public data. Its recommendations are concrete: standardise and unify charging protocols, give cybersecurity equal weight with physical security, upgrade legacy vehicles and stations, design for resilience and interoperability, and raise awareness among drivers and station staff [1]. These recommendations are still valid.

Five limitations need to be borne in mind. First, the paper dates from 2020; it notes that no detailed public security assessment of OCPP 2.0.1 was available then, and it comes before the work on charger management systems and mobile applications covered in Section 5. Second, the grid-impact figures come from simulations that rest on stated

assumptions, such as an omniscient attacker or artificial test networks, so they show plausibility rather than likelihood. Third, defences are treated qualitatively; the paper offers no detection method and no performance figures. Fourth, the attack tree contains no probabilities or costs, so its branches cannot be ranked. Fifth, the paper points out that the proposed security measures remain recommendations that have been neither standardised nor enforced, so its conclusions about deployment depend on regulatory developments that must be checked separately.

4. CRITICAL ANALYSIS OF THE REFERENCE PAPERS

Table -3 summarises the five supplied papers. Three are surveys and two are detection studies that include experiments. The rest of this section appraises each of them in turn.

4.1 Systematic review of charging infrastructure security

Pawlik et al. [2] followed a protocol modelled on PRISMA. A Scopus search returned 196 records, a keyword filter eliminated 94, and 102 publications were analysed. The authors count 26 studies for 2017–2020 and 76 for 2021–2024, a rise of about 192 percent. Authentication and cryptography is the most heavily studied domain, charging infrastructure is the most frequent target, and blockchain is a smaller but growing theme. The review recommends hardening ISO 15118 stacks and certificate workflows, privacy-preserving authentication and billing, and grid-aware intrusion detection adapted to EV-induced oscillations, false data injection and price manipulation in vehicle-to-grid (V2G) services. It further calls for field-grade datasets with open telemetry aligned to IEC 62351-7, and lists the cost of edge computing, the vulnerability of backends to distributed denial of service (DDoS), inconsistent standards and immature explainable AI as barriers.

The review is open about its counts, which is a strength. Its weaknesses are that it draws on a single database, which may overlook work indexed elsewhere, and that studies are tallied by theme instead of being weighted by the quality of their evidence. The authors read the lack of statistically significant differences between categories as a sign that the field has matured; an equally plausible explanation is that the categories are too coarse to show differences. This point is the author's interpretation and not a claim made by the review.

Table -3 Summary of the five reference papers

Ref .	Type and scope	Main contribution	Main limitation
[2]	Systematic review; EV charging attacks, detection and authentication; 2017–2024	102 studies drawn from one database; research output rose from 26 to 76 studies between 2017–2020 and 2021–2024	Single database; studies are counted, not weighed by quality or effect size
[3]	Survey; four taxonomies (domains, CIA-based attacks, defences, verification)	204 of 340 screened papers retained; wide coverage of EV, charger and grid topics	Descriptive; no quantitative comparison; about 25 percent of cited studies give no tools
[4]	Survey; AI-based intrusion detection for in-vehicle networks, 2016–2022	Taxonomy of 102 proposals by attack, frame and feature; algorithm; datasets and AI-model security	Limited to the vehicle bus; no charger or grid link
[5]	Experiment; LSTM autoencoder applied to OCPP traffic indicators from a 12-station site	Modular monitoring platform; unsupervised detection of flooding and disconnections	Only 22 attack samples in the DoS test; headline metrics are class averages
[6]	Experiment; ANFIS classifiers applied to CAN frames from a public dataset	Binary and multiclass detection using SMOTE; 99.6 percent accuracy	Labels taken from an input feature; tuning done on the test set

4.2 Four-way taxonomy of EV cyberattacks

Aljohani and Almutairi [3] screened 340 papers from 2008–2024 drawn from six academic databases and retained 204. They arrange the field into research domains (charging infrastructure, communication, powertrains, power electronics, data privacy, standards), threats classified by confidentiality, integrity and availability, defences, and verification methods. The verification taxonomy is especially useful: it reveals that the studies depend mostly on MATLAB, hardware-in-the-loop (HIL) setups and Python, and that about a quarter of the surveyed works name no tools at all.

The survey is broad but largely descriptive. It states individual results, such as detection rates or suppression

rates, without bringing them onto a common footing. The text also contains some inconsistency: the charging section says the physical layer is not itself open to cyberattacks, yet the power-electronics section later describes electromagnetic interference attacks on sensors and actuators. The distinction can be defended if cyberattacks are restricted to network paths, but this should be said explicitly.

4.3 AI-based intrusion detection for in-vehicle networks

Rajapaksha et al. [4] survey 102 AI-based intrusion detection proposals for in-vehicle networks published from 2016 to August 2022, with the CAN bus as the principal target. Their taxonomy combines attack type, CAN frame features and supervised or unsupervised algorithms, and they additionally review benchmark datasets and the security of the AI models themselves. They conclude that unsupervised methods such as one-class support vector machines and autoencoders produce promising results, that deep learning generally performs better than traditional machine learning, and that LSTM and GRU ensembles and hybrid rule-plus-learning designs can overcome many limitations. Transfer, one-shot, zero-shot and federated learning are named as future trends. Only two of the reviewed proposals deal with protecting the model itself against attack, which the authors regard as a challenging open direction.

For the present study, the survey is valuable because it shows what is known about detection on the vehicle side and how little evidence exists on adversarial robustness. By design, its limit is that it ends at the vehicle boundary and does not consider charger traffic or grid effects.

4.4 LSTM autoencoder detection on charger traffic

Terruggia et al. [5] assembled a monitoring platform from packet capture, a commercial passive sensor, SNMP and syslog agents and an Elasticsearch–Logstash–Kibana stack, and evaluated an LSTM autoencoder at a real site with twelve stations from five manufacturers running OCPP-J 1.6. Nine stations are connected locally and three reach the controller over the public Internet. Fourteen TCP-level indicators are sampled every 5 s. The encoder consists of three LSTM layers with 128, 64 and 32 units, and training on normal traffic used a look-back of 5, a batch size of 128, a learning rate of 10^{-4} and 300 epochs. Anomalies are flagged when the reconstruction error rises above a threshold, which was set at 0.166 after a sweep. The anomalous test data came from simulated flooding with ICMP, SYN and RST packets, sent on 16 July from a host in the same local network to one dedicated station. The model was also tested on station disconnection events.

The abstract of the paper reports 97.1 percent accuracy, 98.6 percent recall and 52 percent precision. The confusion matrix printed for the flooding test holds 17,366 true negatives, 516 false positives, 0 false negatives and 22 true positives. Using standard definitions yields an accuracy of 97.1 percent, which agrees. Precision on the attack class, however, is $22/(22+516)$, or 4.1 percent, and recall on that class is 100 percent. The figures of 52 and 98.6 percent are reproduced exactly by averaging the two classes: class-averaged precision is $(4.1+100)/2 = 52.0$ percent and class-averaged recall is $(100+97.1)/2 = 98.6$ percent. The body of the paper does give the attack-class precision of about 4 percent and an F1 score of about 8 percent, but the abstract and the conclusions present the averaged values as the detector's performance. A reader who depends on the abstract would overestimate how useful the alarms are. Table -4 and Chart -1 present the recomputed values, and the same check on the disconnection matrix gives 54.5 percent precision, 99.2 percent recall and an F1 of 70.3 percent for the attack class, in agreement with the text.

Four further observations follow. First, the DoS test contains only 22 attack samples, from a single day and a single station, and the attacks were generated by a simplified flooding module running inside the local network; the paper itself remarks that many local networks implicitly trust internal hosts. Generalisation to other attack types, such as replay, man-in-the-middle or data injection, which the authors list as future work, is not tested. Second, the threshold was selected by evaluating candidate values on the test data, and the text does not describe a separate validation set containing anomalies; if the same data guided both the choice and the evaluation, the reported operating point is optimistic. Third, the indicators are rolling counters over 5, 15 and 30 minutes, which smooth the signal; this fits the authors' remark that flagged points extend beyond the attack interval, but it also suggests that the timing of alarms is coarse. Fourth, the 516 false alarms could be a few long episodes or many isolated alerts, and the paper does not report alarm episodes, which is what operators would actually experience. The raw data are available only on request, which limits independent checking. Even so, the work is valuable as one of the few evaluations at a real multi-vendor site, and its frank reporting of the low attack-class precision deserves credit.

4.5 ANFIS-based detection of CAN bus attacks

Al Isawi et al. [6] put forward ANFIS classifiers whose initial fuzzy systems are produced by subtractive clustering (ANFIS-SC) or fuzzy c-means (ANFIS-FCM). They use 20,000 samples from a public intrusion dataset recorded through the OBD-II port of a real vehicle while messages were being injected, divide them 80/20 at random, and balance the training set with SMOTE using 30 neighbours. The inputs are the CAN identifier, the data length code and eight data bytes. The accuracies reported are 99.58 and 99.68 percent (binary) and 99.65 and 99.68 percent

(multiclass), which exceed machine-learning baselines such as linear discriminant analysis (85.93 percent) and random forest (99.28 percent) and deep-learning baselines such as CNN (82.63 percent) and LSTM (99.55 percent). Training times are reported to be shorter than those of the deep networks.

Several points weaken these results. The class label is derived from the CAN identifier (0x000 for flooding, 0x164 for impersonation), and the identifier is also one of the ten inputs, so the task may largely come down to recognising two identifier values; the paper does not analyse this possibility. Hyperparameters were chosen by trial and error while evaluating on the test sets, which can inflate performance. The gap between ANFIS-FCM and the LSTM baseline is 0.13 percentage points, which on a 4,000-sample test set amounts to about five samples, and no repeated splits or confidence intervals are provided. A random split of time-ordered frames may put near-identical neighbours in both the training and test sets. The paper presents the study as a detection method for EV CAN buses, yet the dataset description says only that the data came from a real vehicle; whether it was an EV is not stated. Finally, only two injection-type attacks are covered, while fuzzing and replay attacks are not. The contribution is a sound exploration of neuro-fuzzy classification, but the figures should not be taken as proof of robust detection in an operating EV.

5. ADDITIONAL LITERATURE

Five more papers were chosen to cover the period following the base paper and the attack surfaces that the supplied detection studies leave unaddressed. The descriptions below are restricted to what their abstracts and indexing records state; their full texts should be read before specific figures are cited. Table -4 gives an overview.

Johnson et al. [8] classify publicly disclosed vulnerabilities by interface type and cover conventional AC chargers, DC fast chargers and extreme fast chargers. They also review the effects on the power system and on billing, and they identify established IT and operational technology practices as a basis for protection. This is the closest update to the catalogue in the base paper. Hamdare et al. [9] present a risk analysis of charging networks that distinguishes infrastructure-centric from protocol-centric weaknesses, constructs attack scenarios, and examines real charging-session data to test whether the threats can be seen there. It is useful because it links threat descriptions to operational data, something the base paper does not do.

Table -4 Additional peer-reviewed papers used in this study

Ref .	Focus	Finding as stated in the abstract	Link to the base paper
[8]	Review of charger (EVSE) vulnerabilities, organised by interface	Flaws are present in charger devices, vehicle links and upstream services; impacts range from local to national; IT and OT practices can be applied	Updates the vulnerability catalogue of [1]
[9]	Risk analysis of charging station networks	Infrastructure- and protocol-centric vulnerabilities with attack scenarios, tested against real charging-session data	Adds validation based on data
[7]	Impact of EV-based attacks on grid operation	EV loads have high reactive demand; a 30 MW EV attack destabilised a system that withstood a 48 MW residential attack	Quantifies the concern in [1]
[10]	Security analysis of 16 charger management systems	Firmware, mobile and web components assessed; feasibility of attacks on chargers and grid scenarios demonstrated	Extends the backend branch of the attack tree
[11]	Security of EV charging mobile applications	Static and dynamic analysis; apps do not adequately verify EV ownership	Extends the smartphone entry point

Sayed et al. [7] investigate how compromised EVs can be turned against the grid. Their argument is that charging load is non-linear and draws high reactive power, so a smaller attack using EV loads can cause more damage than a larger attack using ordinary residential loads. In their simulations on two grids, a 30 MW EV attack destabilised the system, whereas the grid recovered from a 48 MW residential attack. They also propose patches for known vulnerabilities and two detection mechanisms designed for EV attacks. This strengthens the base paper's grid argument, with the usual caveat that it rests on simulation; the description given here follows the openly available preprint, and the journal version should be checked for differences.

Nasr et al. [10] analyse 16 charger management systems, covering firmware, mobile and web components, through system lookup, reverse engineering and penetration

testing. They demonstrate that the systems can be attacked remotely and discuss attack scenarios against stations and the power grid. Saredidine et al. [11] examine widely used charging applications with static and dynamic techniques and report inadequate verification of EV ownership. Both papers show that the backend and smartphone branches of the base paper's attack tree are not merely hypothetical, since they can be tested on deployed software. They also reveal a gap in detection research, because neither the supplied detection studies nor the surveys describe detectors that watch for abuse of management interfaces or application programming interfaces.

6. RESEARCH GAP

6.1 Where the evidence is concentrated

Table -5 compares how mature the evidence is at each layer. The ratings are the author's qualitative judgement based on the eleven sources of this study and are not a statement about the literature as a whole.

Table -5 Maturity of evidence by layer (author's qualitative rating within the eleven sources)

Layer	Threat evidence	Detection evidence	Grid-consequence evidence
Vehicle bus	Strong [1], [4]	Many offline studies [4], [6]; dataset-bound	Weak; no link to charging load
Charger and network	Strong [1], [8], [9]	Limited; one real site [5]	Moderate; simulated [1], [7]
Backend and apps	Strong, recent [10], [11]	Weak in the sources studied	Moderate; scenarios [10]
Power grid	Strong modelling [1], [7]	Weak; two mechanisms proposed [7]	Simulation only

The pattern is plain. Threat evidence is abundant at every layer, detection evidence is concentrated on the vehicle bus and on a single charger site, and consequence evidence at grid level rests on simulation. None of the sources studied here measures how a charger-level or vehicle-level alarm relates to the quantities that matter to a grid operator, such as the share of demand that has to be disturbed before protection acts. The estimates in [1] and [7] give such thresholds in simulation, but the detection studies [5], [6] do not report time-to-detect or the proportion of stations covered.

6.2 Reliability of reported performance

Two features recur in the two detection studies. Both report very high accuracy on data in which attacks are rare or trivially separable: in [5] the attack class makes up 0.12 percent of the flooding test, and in [6] the label follows an input field. Under such conditions accuracy is a poor guide. Attack-class precision and the precision-recall curve are the informative measures; [5] itself reports an average precision of 0.13 even though the area under the ROC curve is 0.995, which is the typical signature of a rare-event problem. Future studies should report attack-class precision and recall, the number of alarm episodes per day, and time-to-detect, and should select thresholds on a validation set kept separate from the test set.

6.3 Cross-cutting issues

Legacy and migration. The base paper urges the upgrading of legacy vehicles and stations [1], and the 2025 field study still operates OCPP-J 1.6 [5]. Security profiles are available in later protocol versions [3], but their adoption in the field and their assessment remain open. Data. Detection studies depend on single testbeds, simulated environments or public datasets [2], [5], [6]; the review in [2] explicitly asks for field-grade datasets. Robustness. The security of the models themselves is seldom addressed [4]. Operations. False alarms, the cost of edge computing and explainability are identified as barriers [2], and the 2025 field study illustrates the false-alarm problem concretely [5]. Verification. The survey in [3] finds that roughly a quarter of studies do not name their tools, which makes reproduction difficult.

7. CONCEPTUAL ARCHITECTURE

The gaps identified in Section 6 suggest that detection should be arranged in layers and correlated, since no single vantage point sees the full attack path of Fig -1. Fig -2 outlines a conceptual architecture. At the vehicle, a CAN intrusion detector of the types reviewed in [4] would monitor in-vehicle traffic. At the charger, a sequence-aware detector would use traffic indicators similar to those in [5]. At the backend, audit logs of management and application interfaces would be examined for abuse patterns like those documented in [10], [11]. On the physical side, station power quality and setpoint consistency would be monitored, in line with the converter-level effects reported in [1]. A correlation layer would search for fleet-wide synchrony and for inconsistencies between commanded and observed aggregate demand, which are the signatures of the two attack-tree branches in [1]. Its output would drive a grid-aware response such as isolating a station, rate-limiting commands and alerting the operator and utility.




Fig -2: Proposed the Conceptual layered monitoring architecture

8. LIMITATIONS OF THIS STUDY

This is a narrative synthesis of eleven sources and not a systematic review, so its coverage of the wider literature is partial. The descriptions of the five additional papers rely on abstracts and indexing records, which is why detailed claims about them are deliberately restricted. The qualitative ratings in Table -5 reflect the sources selected and involve judgement. The recomputation in Section 4.4 depends on the confusion matrices as printed in the original paper; if the authors used a different evaluation set for the headline figures, the discrepancy may have an explanation that cannot be seen in the paper. The architecture in Section 7 is untested. The literature was reviewed as of October 2026, and newer work may exist.

9. CONCLUSIONS

The base paper [1] remains a sound framework: the chain from vehicle to station to grid is real, the interfaces are many, and the logic of the attack tree is reasonable. Later work confirms the vulnerability side, including management systems [10] and smartphone applications [11], and backs the grid concern with simulation evidence [7]. The detection side is less mature than its headline numbers imply. In the 2025 charger-traffic study [5], the labelling scheme may make the task easier than it seems. Neither study links its detector to grid-level consequences.

The most valuable next steps are methodological and architectural: assess detectors against grid-level thresholds in co-simulation, extend detection to backend and application interfaces, and publish field-grade data. Until such evidence is available, claims that AI-based detection protects the charging-to-grid chain should be regarded as promising hypotheses and not as established results.

REFERENCES

- 1) S. Acharya, Y. Dvorkin, H. Pandžić, and R. Karri, "Cybersecurity of smart electric vehicle charging: A power grid perspective," *IEEE Access*, vol. 8, pp. 214434–214453, 2020, doi:10.1109/ACCESS.2020.3041074.
- 2) L. Pawlik, J. L. Wilk-Jakubowski, P. T. Grabski, and G. Wilk-Jakubowski, "Securing the electrified future: A systematic review of cyber attacks, intrusion and anomaly detection, and authentication in electric vehicle charging infrastructure," *Energies*, vol. 18, no. 18, Art. no. 4847, Sep. 2025, doi:10.3390/en18184847.
- 3) T. Aljohani and A. Almutairi, "A comprehensive survey of cyberattacks on EVs: Research domains, attacks, defensive mechanisms, and verification methods," *Defence Technology*, vol. 42, pp. 31–58, 2024, doi:10.1016/j.dt.2024.06.009.
- 4) S. Rajapaksha, H. Kalutarage, M. O. Al-Kadri, A. Petrovski, G. Madzudzo, and M. Cheah, "AI-based intrusion detection systems for in-vehicle networks: A survey," *ACM Comput. Surv.*, vol. 55, no. 11, Art. 237, Feb. 2023, doi:10.1145/3570954.
- 5) R. Terruggia, A. Maldarella, G. Dondossola, and G. Webber, "Enhancing the detection of cyber-attacks to EV charging infrastructures through AI technologies," *Electronics*, vol. 14, no. 21, Art. no. 4321, Nov. 2025, doi:10.3390/electronics14214321.
- 6) O. A. Al Isawi, K. A. Al Jaafari, and A. S. Al-Sumaiti, "Electric vehicles CAN bus cyber attacks detection using adaptive neuro fuzzy inference system," *IEEE Access*, vol. 13, pp. 90862–90874, 2025, doi:10.1109/ACCESS.2025.3550970.
- 7) M. A. Sayed, R. Atallah, C. Assi, and M. Debbabi, "Electric vehicle attack impact on power grid operation," *Int. J. Electr. Power Energy Syst.*, vol. 137, Art. no. 107784, May 2022.
- 8) J. Johnson, T. Berg, B. Anderson, and B. Wright, "Review of electric vehicle charger cybersecurity vulnerabilities, potential impacts, and defenses," *Energies*, vol. 15, no. 11, Art. no. 3931, May 2022, doi:10.3390/en15113931.
- 9) S. Hamdare et al., "Cybersecurity risk analysis of electric vehicles charging stations," *Sensors*, vol. 23, no. 15, Art. no. 6716, Jul. 2023, doi:10.3390/s23156716.
- 10) T. Nasr, S. Torabi, E. Bou-Harb, C. Fachkha, and C. Assi, "Power jacking your station: In-depth security analysis of electric vehicle charging station management systems," *Comput. Secur.*, vol. 112, Art. no. 102511, 2022.
- 11) K. Sarriddine, M. A. Sayed, S. Torabi, R. Atallah, and C. Assi, "Investigating the security of EV charging mobile applications as an attack surface," *ACM Trans. Cyber-Phys. Syst.*, vol. 7, no. 4, Art. 26, pp. 26:1–26:28, 2023, doi:10.1145/3609508.