IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429
Volume 6, Issue 6, October 2026
Impact Factor: 8.2

AI-Driven Safety And Thermal Warning for EV Charging

Mansi Santosh Karad[1], Dr Sunil Magan More[2], Arpita Dhananjay Shah[3], Sayma Ashpak Mansuri[4], Deepak Sanjay Sapkal[5], Arya Deepak Korade[6], Santoshi Dnyaneshwar Shinde[7], Siddhi Prashant Warule[8]

[1,3,5,6,7.8] Students, Dept. of Electrical Engineering,
[2] Professor, Dept. of Electrical Engineering
Guru Gobind Singh College of Engineering and Research Centre, Nashik, Maharashtra, India

Abstract: *Charging an electric vehicle (EV) has become an everyday activity, yet a safe session requires far more than the delivery of electrical power. Abnormal voltage or current behavior, rising temperatures, damaged connectors, battery faults, smoke or flame, and cyberattacks that tamper with measurements or control messages can all compromise a session. This paper reviews published research on artificial-intelligence (AI) methods for EV charging safety and builds from it a combined conceptual framework for early warning. The framework joins time-series prediction with an adaptive long short-term memory (A-LSTM) network, anomaly detection with an LSTM autoencoder and other machine-learning models, computer-vision recognition of smoke and fire, and a cyber-intrusion monitoring layer. Sensor readings such as voltage, current, temperature, and state of charge are compared with both predicted behavior and fixed safety rules, while camera evidence and network anomalies provide further inputs to the risk decision. When the risk passes a validated threshold, the system may notify an operator, lower the charging power, or request a controlled shutdown through the existing protection hardware. The studies reviewed report encouraging results for their individual tasks: 98.5% anomaly-detection accuracy in a simulated charging-port-current study, 94.09% mean average precision (mAP) for an improved fire-detection model, and strong performance in time-series charging warning and cyberattack detection. Because these figures originate from different datasets and evaluation settings, they must not be compared directly. The proposed integration is a design framework rather than an experimentally validated product. Real-world deployment would call for charging-station trials, calibration to different vehicles, careful control of false alarms and independent safety validation*

Keywords: Electric vehicles; EV charging safety; thermal warning; artificial intelligence; A-LSTM; LSTM autoencoder; computer vision; thermal runaway; anomaly detection; cyber-physical security

I. INTRODUCTION

Electric vehicles play a central role in the shift to lower-emission transport, and the growing EV fleet is driving the expansion of both public and private charging stations. A modern charging station is a cyber-physical system: besides delivering electrical energy, it exchanges data between the vehicle, the charger controller, a management platform and, in some cases, the wider power grid. This integration makes charging convenient, but it also means that a fault may be electrical, thermal, software-related or tied to communication security [3, 5, 6].

Safety during charging deserves particular attention because a lithium-ion battery packs a great deal of energy into a small volume. Abnormal charging conditions, cell damage, internal faults or inadequate thermal management can cause overheating and, in the worst cases, thermal runaway. As such an event develops, shifts in voltage, current or

IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429
Volume 6, Issue 6, October 2026
Impact Factor: 8.2

temperature may show up before any smoke or flame can be seen. An effective warning system should therefore pick up suspicious trends early instead of waiting for a fire to become obvious [1, 2].

Conventional protection devices remain indispensable. Fuses, contactors, insulation monitoring, battery-management systems and charger interlocks offer direct protection, and AI should complement them by detecting patterns that a single fixed threshold cannot capture. A given voltage, for instance, may be perfectly normal for one vehicle at a certain point in charging yet abnormal for another. Time-series models can learn the expected trend from past data, and computer vision can contribute evidence from visible smoke or flames [1, 2].

A second concern is whether the charging data can be trusted. Spoofed current values, modified messages, denial-of-charging attacks and man-in-the-middle attacks can disrupt service or conceal a hazardous condition, and a safety monitor that depends on one data stream can be deceived if that stream is compromised. The cybersecurity studies reviewed here therefore point towards a layered approach that pairs sensor behaviour with network-level anomaly detection [3, 4, 5, 6].

1.1 Aim and objectives

This paper aims to propose an integrated AI framework for EV charging safety and thermal warning, grounded in the reference studies cited in this paper. Its objectives are to:

- Survey the AI techniques applied to charging-voltage prediction, current anomaly detection, fire and smoke recognition, and cyberattack detection.
- Explain how sensor data, past charging patterns and camera images can together support early warning.
- Put forward a modular workflow that issues graded warnings while leaving conventional protection hardware in charge.
- Summarise the results reported by the reference studies, making clear that they are not new experiments.

Point out practical limitations and set out a realistic plan for future validation.

1.2 Scope of this paper

his is a literature-based design paper. It did not involve collecting a new EV charging dataset, connecting a battery to a test bench or training a new model. All reported numbers are credited to the studies in which they originated. Diagrams described as proposed or conceptual are original explanatory figures, whereas the photographs and research figures are cropped excerpts from the cited papers and are credited in their captions.

II. RELATED WORK

2.1 Current anomaly detection

Mitikiri, Srinivas and Pal investigated adversarial cyberattacks on EV charging stations [3]. Working with a simulated DC charging infrastructure, they used an LSTM autoencoder to capture the time-dependent behaviour of the charging-port current, and a Kolmogorov–Smirnov test compared the distributions of predicted and observed values across sliding windows. The accuracy they reported was 98.5% in their simulated setting. The result illustrates how sequence modelling can help expose spoofed or unusual current data, though it does not guarantee the same accuracy at every real charging station.

IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429
Volume 6, Issue 6, October 2026
Impact Factor: 8.2

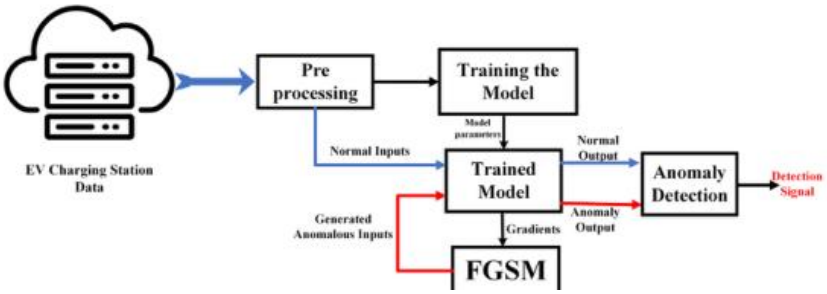

Figure 1. Workflow for generating and detecting anomalous charging-current data, as presented in the reference study. Cropped from Mitikiri et al. [3], p. 3.

2.2 Vision-based fire and smoke detection

Gao and co-authors studied computer-vision approaches to detecting fire while an EV is charging [2]. Their model is an improved YOLOv5 variant in which the feature-extraction stage and an attention mechanism have been modified so that it can pick out visual cues such as flames and smoke. In their test environment it achieved an overall mean average precision (mAP) of 94.09%, an average F1 value of 90.00% and a detection speed of 42.09 frames per second. A camera can monitor a larger area than a single point sensor, but changes in lighting, occlusion, weather and unrelated smoke may trigger false alarms.

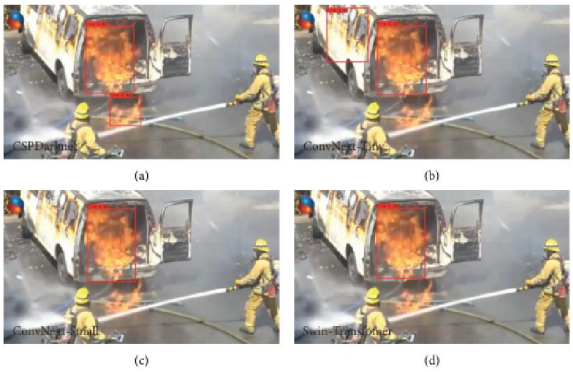

Figure 2. Example fire and smoke detection results, reproduced as a cropped figure from Gao et al. [2], p. 9. The image is included to give context from the literature and does not represent a new experiment.

2.3 AI anomaly detection and load forecasting

Ghoreishi and co-authors paired anomaly detection with short-term forecasting of EV charging load [5]. Autoencoders, isolation forest, one-class SVM and DBSCAN were tested on real session data from charging stations, and the forecasting stage combined SARIMA with a random-forest correction of its residuals. On the filtered dataset, the hybrid forecast reached a mean absolute error of 11.95 kWh, against 13.08 kWh for SARIMA alone and 12.12 kWh for random forest. These findings underline that abnormal sessions should be dealt with before a demand forecast is trained or used.

2.4 Cybersecurity and federated learning

The cybersecurity studies reviewed here investigate both supervised and unsupervised learning against attacks on EV charging networks. EVDetect offers an orchestrated framework in which models are compared on detection quality as

IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429
Volume 6, Issue 6, October 2026
Impact Factor: 8.2

well as computational cost [6]. A separate transfer-learning study applies a CNN–BiLSTM model with progressive domain refinement to carry knowledge from broader security datasets over to EV charging traffic [7]. Biswas and co-authors pair energy-timing watermarking with federated learning, allowing charging stations to improve detection jointly without sharing raw telemetry [4]. These approaches matter here because a safety warning can be trusted only if the measurements and control messages behind it are themselves reasonably trustworthy.

Figure 3. Energy-timing watermarking and federated anomaly-detection architecture from the reference study. Cropped from Biswas et al. [4], p. 7.

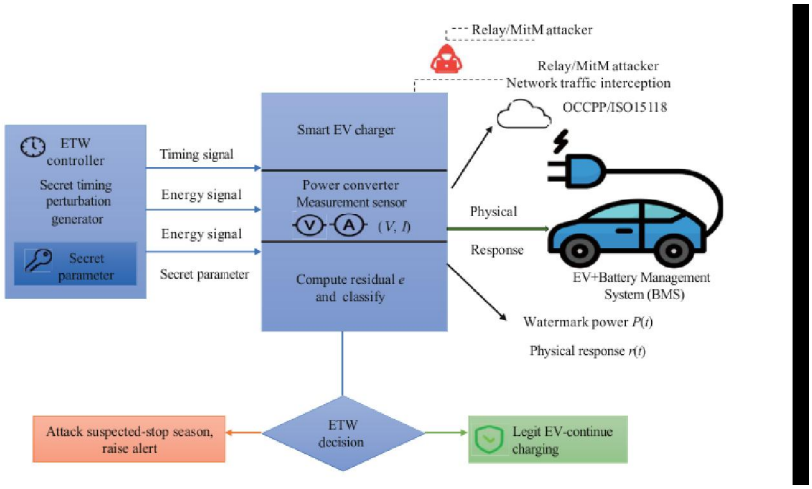

III. PROPOSED INTEGRATED FRAMEWORK

The proposed system consists of several cooperating modules. It is not intended to replace the vehicle's battery-management system (BMS), the charger's protection circuits, the emergency stop or fire-safety procedures. Rather, it adds an intelligence layer that monitors the data, predicts expected behaviour, and recommends or initiates actions through approved control interfaces.

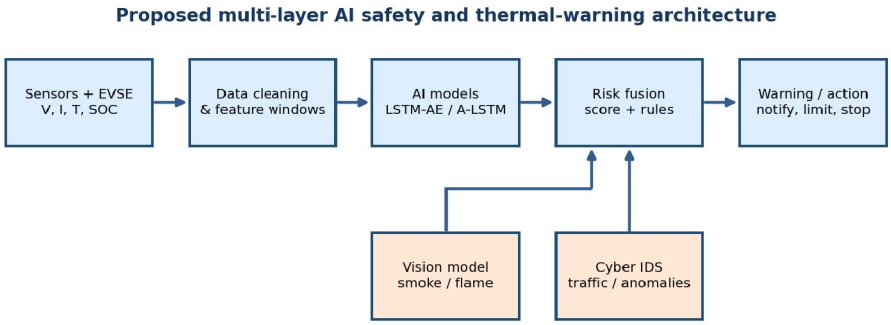

Figure 4. Multi-layer architecture proposed in this paper. The diagram is conceptual and does not represent a validated implementation.

3.1 Data acquisition

The data layer gathers information that is already available, or that can be measured safely, at the charging point. Which signals can be used depends on the vehicle and the charger, but the most useful ones are voltage, current, charging power, connector and ambient temperature, charging duration, state of charge (SOC), error codes, contactor state and communication events. Camera frames can be analysed locally to detect smoke or flame. Timestamps and

IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429 Volume 6, Issue 6, October 2026 Impact Factor: 8.2

sensor-quality flags should be recorded so that missing, delayed or unreliable measurements are never silently treated as normal.

3.2 Pre-processing and feature extraction

Before any prediction is made, the measurements should be screened for sensor dropouts, missing values, out-of-range readings and duplicate timestamps. Parameters obtained from the training set can be used to scale numeric values, and time-series models can be fed sliding windows, such as the latest N measurements together with their recent changes. Useful derived features include the temperature rise per minute, the voltage prediction residual, current variation, power change, the rate of change of SOC and the differences between related sensors.

3.3 Prediction and anomaly models

Using recent observations and historical charging patterns, an A-LSTM model can forecast how voltage or temperature is expected to evolve next [1]. An LSTM autoencoder can learn a compact representation of normal current sequences and flag windows whose reconstruction error is large [3]. Simpler models, such as isolation forest or one-class SVM, can serve as complementary detectors when the data are tabular or the deployment hardware is limited [5, 6]. In a practical design, the models should be compared on the same dataset, and latency and memory use should be assessed alongside accuracy.

3.4 Vision and cyber monitoring

A camera-based model can search for smoke or flame patterns independently of the other channels [2]. In parallel, a cybersecurity module can examine communication features, power-consumption traces and event logs for suspicious activity [3, 4, 6, 7]. The channels should be kept apart in the early stages of processing so that a failed camera or network connection does not disable the sensor-based safety function. Their outputs can then be merged into a risk estimate using transparent rules.

3.5 Risk fusion and response

A basic fusion layer can determine a risk level by weighing several pieces of evidence together. A high temperature residual accompanied by an unusual current pattern, for example, deserves more attention than a single noisy sensor sample. A camera-based fire alert may, according to site policy, trigger an immediate high-priority response. The final response should be

designed together with safety engineers and may consist of logging the event, alerting the driver or operator, requesting a controlled reduction of charging power, or asking the certified charger protection logic to end the session. AI predictions alone should not be allowed to bypass hard safety limits.

IV. METHODOLOGY FOR THERMAL WARNING

4.1 Time-series windowing

Let x(t) denote a sensor value at time t, for example the connector temperature or the charging voltage. The model takes as input a window of the most recent measurements:

$$X_t = [x(t-N+1), x(t-N+2), \ldots, x(t)]$$

where N is the number of samples in the input window. From this window the model estimates the next value, or a short future sequence, denoted $\hat{x}(t+1)$. The residual is the absolute difference between the observed and predicted values:

$$e(t) = |x(t) - \hat{x}(t)|$$

A large residual indicates that the observation departs from the pattern the model expected. A residual on its own, however, does not prove that a dangerous fault exists. The system should also take into account absolute limits, sensor reliability, the stage of charging and whether the deviation persists over several samples.

4.2 Dynamic warning threshold

A dynamic threshold can be computed from recent residuals. A simple illustrative form is:

$$T_t = \mu_e + k\sigma_e$$

IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429
Volume 6, Issue 6, October 2026
Impact Factor: 8.2

Here μ_e and σ_e are the mean and standard deviation of the residuals over a recent window of normal data, and k sets how conservative the warning is. The equation is a general design example and is not a threshold taken directly from any one reference paper. A real system would have to be calibrated for the vehicle, the sensor accuracy, the charging mode, the ambient conditions and the safety requirements. The threshold must also be prevented from drifting so far that a persistent fault comes to be regarded as normal.

Figure 6. Voltage prediction results from the A-LSTM charging- warning study, cropped from Diao et al. [1],p.9.

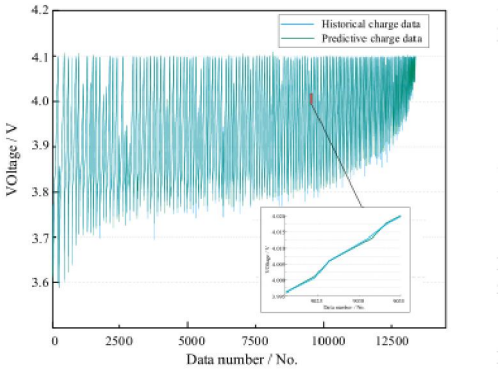

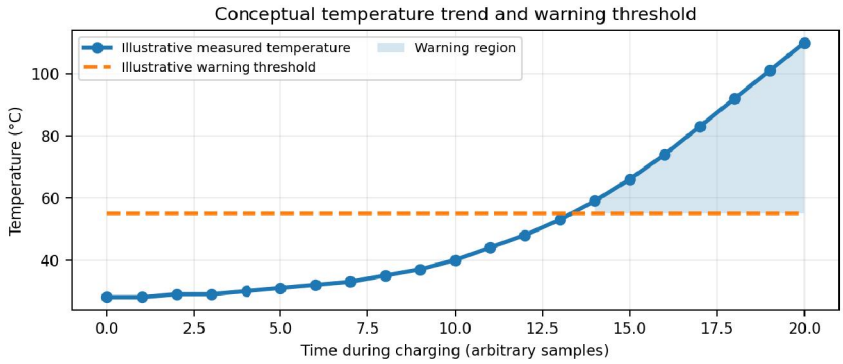

Figure 6. Conceptual illustration of a temperature rise and a warning threshold. The values are for illustration only and are not measured EV data.

4.3 Temperature-specific indicators

Temperature should be judged by its level as well as its rate of change. A connector that warms gradually during a high-current session is a different case from one whose temperature jumps suddenly while the current stays steady. Useful features include T(t), ΔT over a fixed interval, dT/dt, the difference between ambient and connector temperature, and the relationship between current and temperature. Where battery temperature can be obtained through the BMS, it should be used in line with the vehicle manufacturer's interface and data-quality rules. Because a public charger may have no direct access to internal cell temperatures, the design must clearly separate measured values from estimated ones.

4.4 Warning states

Normal: sensor values stay within expected patterns, and no independent detector reports a serious event.

Advisory: a single model detects a mild deviation; monitoring continues and the event is recorded.

Warning: the deviation persists, grows rapidly or is confirmed by more than one independent signal; the operator is notified and, if available, a validated power-limiting policy is applied.

IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429 Volume 6, Issue 6, October 2026 Impact Factor: 8.2

Critical: a hard protection limit is reached, or there is a strong fire or smoke indication or a combination of severe indicators; the site's emergency shutdown and response procedure is followed.

The states above form a proposed, human-readable policy. The exact limits, timing and control actions must be established through engineering validation and the relevant standards rather than chosen arbitrarily for a conceptual model.

V. CYBER-PHYSICAL AND VISION-BASED SAFETY

5.1 Current-sequence anomaly detection

The approach in [3], based on an LSTM autoencoder, suits the problem because charging current is a sequential signal. Once trained on normal data, the model reconstructs each input sequence, and a window whose reconstruction error is unusually large can be flagged for review. The study also applies a statistical comparison between the predicted and observed distributions in sliding windows, which can reveal spoofed data even when every individual value looks plausible. Before deployment, the model should be tested against harmless variation, for example between EV models, charge levels, ambient temperatures and charging rates.

5.2 Camera-based smoke and fire detection

Computer vision contributes a different type of evidence: it can spot visible smoke or flames near a vehicle or a charging pile. The improved YOLOv5 method in [2] aims to detect fire-related visual features better while still running in real time. A camera cannot replace temperature and electrical sensors. Its performance can be degraded by rain, glare, darkness, steam, dust or an obstructed lens, so camera confidence should be combined with health checks and with other signals.

5.3 Network-level intrusion detection

The cybersecurity papers reviewed here address attacks including man-in-the-middle manipulation, spoofing, false data injection and denial of service. Supervised models are able to classify known categories of attack, whereas unsupervised detectors can highlight unfamiliar patterns for follow-up investigation [3, 4, 6]. Transfer learning may lessen the need for EV-specific labelled data, but performance can change when the charger, protocol, firmware or type of attack differs from the training domain [7]. The system should therefore report its uncertainty and retain logs for later analysis.

5.4 Privacy-aware learning

Charging records can expose patterns of vehicle use and user routines. Federated learning is one way to improve models across several stations without gathering all raw telemetry in one place. In the watermarking approach described in [4], each local station shares model updates but keeps its raw measurements locally. Federated learning does not by itself guarantee privacy or security; the protection of updates, station authentication, data governance and the risk of model poisoning all still require attention.

5.5 Combined decision logic

A practical decision layer can combine a rule-based structure with model scores. A high temperature together with a rapid temperature rise, for instance, may be classed as a thermal concern; a large current residual coinciding with a suspicious communication event may point to a cyber-physical problem; and a high-confidence smoke or fire detection can escalate the event on its own. The combined logic should be designed to fail safely when a sensor, camera or network connection becomes unavailable.

VI. EVIDENCE FROM THE REFERENCE PAPERS

The reviewed studies address different tasks and use different datasets and environments. The table below summarises selected values as the authors reported them. They provide useful evidence that AI is able to support several parts of the proposed framework, but they should not be ranked as though all the studies had used the same test set.

IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429
Volume 6, Issue 6, October 2026
Impact Factor: 8.2

Reference / task	Method	Reported result	Important context
Diao et al. [1] Charging warning	Adaptive LSTM; dynamic voltage threshold	Average warning accuracy of 99.0% reported for selected simulated fault types	Based on daily charging data with simulated fault cases and dynamic thresholds; not a general fire detector.
Gao et al. [2] Fire/smoke vision	Improved YOLOv5 with attention	mAP 94.09%; F1 90.00%; 42.09 FPS	Computer-vision benchmark obtained under the paper's own test setup.
Mitikiri et al. [3] Current anomaly	LSTM autoencoder + KS test	98.5% accuracy	Simulated DC charging setup with adversarial inputs.
Ghoreishi et al. [5] Load forecasting	SARIMA + random-forest residual correction	MAE 11.95 kWh on filtered data	Forecasting error; not a safety-detection accuracy.
Lazaridis et al. [6] Cyber intrusion	EVDetect model orchestration and ML classifiers	F1-macro of about 0.87–0.88 in the full binary experiment	Dataset and model-comparison setting; deployment on real electric vehicle supply equipment (EVSE) remains to be validated.
Kesavan et al. [7] Transfer-learning IDS	CNN–BiLSTM with domain refinement	Accuracy of 98.25%; mean latency of 18.4 ms per sample reported	Simulated or emulated attacks; the authors note that real-world validation is left for future work.
Biswas et al. [4] MitM detection	Energy-timing watermark + federated learning	Accuracy of about 99.7% and F1 of 99.85% reported	IoTID20 dataset used as a proxy, with synthetic features resembling energy-timing watermarks (ETW); not a physical EV battery experiment.

Copyright to IJARSCT
www.ijarsct.co.in
DOI: 10.48175/568

IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429
Volume 6, Issue 6, October 2026
Impact Factor: 8.2

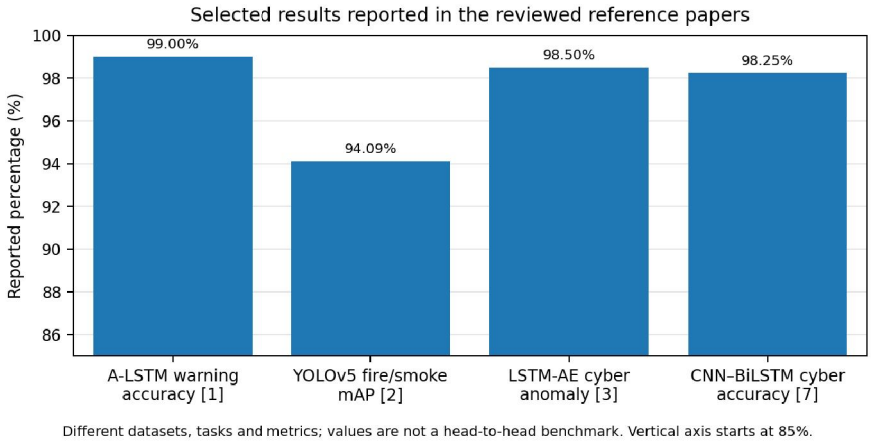

Figure 7. Selected metrics reported in the reviewed papers. The metrics belong to different tasks and cannot be compared directly; the rounded A-LSTM value is shown only for visual orientation.

6.1 Interpretation

The literature indicates that time-series prediction can reveal deviations before a fixed threshold is reached, that computer vision can recognise visible signs of fire, and that cybersecurity models can identify unusual charging behaviour or suspicious network traffic. The most important practical lesson is not that one model can solve every problem, but that detectors operating separately can cover different failure modes. Reported performance may drop when the data differ from the training conditions, so validation should cover different chargers, vehicles, seasons, lighting conditions and fault types.

VII. PROPOSED EVALUATION PLAN

Any future implementation should be tested in stages, starting with offline data and progressing to a controlled pilot. Testing must not deliberately expose people or vehicles to unsafe charging conditions, and fault-injection and adversarial tests should be carried out in simulation or on a properly isolated test setup.

Dataset definition: assemble a labelled dataset of time-stamped voltage, current, temperature, SOC, charging-state, alarm and camera-event records, with data-quality information attached.

Train/test separation: divide the data by charging session and, where possible, by vehicle or charging station, to limit data leakage.

Baseline comparison: benchmark the dynamic threshold, A-LSTM, LSTM autoencoder and isolation forest against suitable vision and intrusion-detection system (IDS) baselines.

Metrics: report precision, recall and F1-score together with the false-alarm rate, missed-event rate, detection lead time, inference latency, and memory and CPU usage.

Scenario testing: include sensor failure, abnormal heating, current spoofing, communication loss and camera outages, as well as normal charging variation and smoke-like visual confounders.

Safety review: confirm that AI output can never override certified protection limits, and that a model crash or loss of communication drives the system to a defined safe state.

Pilot evaluation: before any wider deployment, measure performance across several charging points and under varying ambient conditions.

7.1 Metrics

In a binary warning task, precision is TP/(TP+FP) and recall is TP/(TP+FN), with TP denoting a correctly detected hazardous event, FP a false alarm and FN a missed event. Accuracy by itself can be misleading when hazardous events are rare. A system that overlooks a real thermal event may be more dangerous than one that raises an occasional false

IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429
Volume 6, Issue 6, October 2026
Impact Factor: 8.2

alarm, yet too many false alarms can also cause operators to disregard warnings. Detection lead time is particularly important, since a correct alert that comes too late may offer little safety benefit.

7.2 Deployment considerations

The model must be compact enough for the target hardware and quick enough to keep up with the sampling rate. Immediate warnings can come from a local edge computer at the charger, while a server handles longer-term analysis. Sensors, video and network logs need to be time-synchronised before they can be compared. Model versions should be tracked, and every change should be tested before deployment. Operators should be given an understandable explanation, such as 'connector temperature rising faster than expected', and not merely a model score.

VIII. LIMITATIONS AND CHALLENGES

The principal limitation of this paper is that it derives an integrated design from existing studies and does not report a new controlled experiment. The reference papers rely on different datasets and assumptions. Several of their results rest on simulated charging data, proxy cybersecurity datasets or emulated attacks, and none of them yields a single end-to-end safety rate for the combined system [3, 4, 6, 7].

Thermal runaway is a complex phenomenon and does not always give a clear, early external warning sign. A temperature sensor may sit far from the hottest point, and the charger may have no access to the internal battery temperature. Camera systems can be degraded by environmental conditions. AI models may also absorb biases from incomplete datasets, drift over time, or produce false alarms when a new vehicle behaves differently from the training examples.

Cybersecurity poses a further challenge, because an attacker may target the sensors, the inputs to the models, the communication channel or the model itself. A system built from several models can lessen the dependence on any one signal, but it may still fail if several channels share the same vulnerability. Federated learning can support distributed training, yet it demands secure aggregation and careful handling of model updates [4].

IX. DISCUSSION

An integrated AI warning system delivers the most value when it supports, rather than replaces, established engineering safeguards. The time-series layer searches for abnormal trends in electrical and thermal measurements, the vision layer for visible evidence of smoke or fire, and the cyber layer checks whether the data and messages can be trusted. The risk-fusion layer then turns this evidence into a graded warning and passes it to approved protection logic.

The design also promotes a clear separation between early prediction and emergency protection. AI is able to estimate the risk and let an operator act sooner, but hard safety limits should stay deterministic and be tested independently. A model might, for example, warn that the temperature is climbing unusually fast before any fixed threshold is reached, while a certified protection system stays in charge of cutting the power once a required limit is exceeded.

Explainability is a further benefit. A warning should state the main reason for the alert together with the signals that support it. A transparent message lets an operator judge whether to inspect a connector, verify a sensor, examine the communication
logs or activate a fire response plan. A deployment should also maintain an audit trail covering model outputs, thresholds, data quality and operator actions.

X. CONCLUSION

This paper reviewed AI-driven safety and thermal warning for EV charging on the basis of the reference studies and presented a conceptual integrated framework. The methods examined demonstrate that computer vision, adaptive LSTM prediction, LSTM autoencoders, classical anomaly detection, federated cybersecurity and transfer learning can each tackle a different part of the problem. The proposed framework brings together sensor prediction, temperature-trend monitoring, visual fire detection and cyber-anomaly analysis, and maps their outputs onto clear warning levels.

IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429
Volume 6, Issue 6, October 2026
Impact Factor: 8.2

The results reported in the reference papers are encouraging, but they were obtained on different datasets and under different testing conditions, and they do not prove that the proposed combined system already attains such values. The logical next step is to build and test the modules in a controlled environment, assess false alarms and missed events, try realistic fault scenarios, and verify every action against the existing electrical and battery safety protections. If validation is careful and the warnings are designed around human users, AI can become a valuable extra layer for safer, more dependable EV charging.

REFERENCES

[1] X. Diao, L. Jiang, T. Gao, L. Zhang, J. Zhang, L. Wang, and Q. Wu, "Research on Electric Vehicle Charging Safety Warning Based on A-LSTM Algorithm," IEEE Access, vol. 11, pp. 55081–55093, 2023. doi: 10.1109/ACCESS.2023.3281552.

[2] Y. Gao, Q. Yang, S. Zhang, and D. Gao, "Real-Time Fire Detection Method Based on Computer Vision for Electric Vehicle Charging Safety Monitoring," Journal of Electrical and Computer Engineering, 2023, Article ID 9215528. doi: 10.1155/2023/9215528.

[3] S. B. Mitikiri, V. L. Srinivas, and M. Pal, "Anomaly detection of adversarial cyber attacks on electric vehicle charging stations," e-Prime – Advances in Electrical Engineering, Electronics and Energy, vol. 11, 100911, 2025. doi: 10.1016/j.prime.2025.100911.

[4] P. Biswas, S. K. Singh, and V. Rishiwal, "An information-driven and intelligent federated energy-timing watermark framework for man-in-the-middle attack detection in electric vehicle charging networks," Journal of Information and Intelligence, vol. 4, pp. 311–327, 2026. doi: 10.1016/j.jiixd.2026.04.005.

[5] M. Ghoreishi, F. Santori, A. Geri, M. Jabari, F. Bellesini, M. Maccioni, and T. Bragatto, "AI-powered anomaly detection and load forecasting for sustainable EV charging networks," Energy Reports, vol. 15, 109186, 2026. doi: 10.1016/j.egyr.2026.109186.

[6] K. Lazaridis, I. Topouzelidis, I. Vagionas, and A. Lekidis, "EVDetect: a framework for orchestrated and AI-powered detection of cyber-attacks in EV charging networks," Scientific Reports, Article in Press, 2026. doi: 10.1038/s41598-026-67199-1.

[7] V. T. Kesavan, R. Gopi, M. J. Hossen, E. R. Joseph, and D. Danalakshmi, "Early detection of cyberattacks in renewable-integrated EV charging systems using transfer learning," Scientific Reports, Article in Press, 2026. doi: 10.1038/s41598-026-71041-z.

Figure source and reuse note

Figures 1–3 and 5 are cropped excerpts from the cited reference papers and are credited to those papers in their captions; the publishers' licence and reuse conditions apply to them, and permission should be confirmed before publication. Figures 4, 6 and 7 are original explanatory illustrations prepared for this paper. The temperature curve in Figure 6 is conceptual, and the chart in Figure 7 is a visual summary of selected published results rather than a new experiment.