5G-enabled Temperature Sensor Fusion & Federated Learning for optimal operation of EV charging points and fault prevention

Alexios Karadimos
Department of Electrical and
Computer Engineering
National Technical University of Athens
Email: akaradimos@mail.ntua.gr

Christos Stefanatos
Department of Electrical and
Computer Engineering
National Technical University of Athens
Email: cstefanatos@epu.ntua.gr

Evanthia Sismanoglou
Department of Electrical and
Computer Engineering
National Technical University of Athens
Email: esismanoglou@mail.ntua.gr

Vaggelis Marinakis
Department of Electrical and
Computer Engineering
National Technical University of Athens
Email: vmarinakis@epu.ntua.gr

Abstract—The increased adoption of Electric Vehicles (EV) necessitates smart and robust EV Supplying Equipment (EVSE) that can reliably charge EVs. Reliable EV charging means offering the right amount of Electric Power given the Power that the Energy grid can give at each time and also based on the Vehicles and Charging Point's characteristics. A key problem that is observed, especially in Mediterranean countries with hot summers, is the high temperature values that are observed during charging. Very often, the maximum temperature is exceeded causing faults in the EVSE hardware, power losses, and interruption of the charging sessions. This paper proposes a smart framework to monitor the temperature of the EV and the EVSE, from a variety of inputs that are then fused together (sensor fusion), and pipelined to a Federated Learning algorithm to avoid hardware faults and suggest the optimal power so that the charging session continues reliable within safe operating temperature limits. The sensor fusion algorithm runs locally at each charging point, the results are then pipelined to a local Machine Learning algorithm and then the model parameters are sent via 5G to a cloud central FL model.

I. INTRODUCTION

A. EV Charging

The last decade, there has been an almost exponential growth of EVs [1], [2] which has as a result a lot of manufacturers to offer AC and DC charging solutions to meet the demand in EV charging infrastructure. A systematic approach to plan charging points based on location, congestion, future demand, and grid constraints is presented in [2]–[4], but these studies do not plan for the risk of charging points being faulted due to high temperatures, causing a risk for potential EV drivers to not be able to charge their vehicles.

B. Sensor Fusion

Sensor fusion is a vital algorithm in modern IoT systems, that combines data from multiple sensors -such as temperature, voltage and current sensors-, allowing for more accurate calculation and monitoring of temperature, voltage or other variables of a system. It is widely used in EV Battery Management Systems (EV-BMS), combining data from multiple sources such as temperature sensors, voltage sensors, and current sensors to improve state-of-charge (SOC) estimation, thermal management, and overall system safety [5]. Advanced sensor fusion techniques may also use other models like long short-term memory (LSTM) and Bayesian networks. For instance, the application of LSTM models in combination with Bayesian networks has shown promising results in predicting EV charging loads and energy consumption patterns [6].

C. Federated Learning

Federated learning is a distributed machine learning approach that allows model training on decentralized data and has gained attention in the EV charging domain due to its ability to leverage data from multiple charging points while preserving privacy. This approach is particularly relevant in the context of 5G-enabled charging infrastructure, where high-speed, low-latency communication facilitates real-time data exchange and model updates across distributed charging stations [7].

D. The problem

All EV chargers have specific recommended temperature ranges, within which Electric Power can be supplied to Electric Vehicles reliable and without power losses or disruptions. However, smaller AC chargers, which usually are constructed with cheaper materials (e.g Plastic), are smaller in size and usually do not have cooling mechanisms, are more susceptible to reach higher temperatures, especially in Mediterranean countries with hot summers.

II. SYSTEM ARCHITECTURE

The proposed framework is organized into two main parts: the sensor fusion and Machine Learning model at the edge level (EVSE) and the central model and cluster managements at the cloud level.

A. Edge Level

At the edge level the proposed system is organized as shown in 1. Data points are collected from internal and external temperature sensors, current and voltage sensors as well as temperature values from weather forecasts and the EV charging point temperature model. All these data are then fused together using a Kalman filter which outputs a more robust and accurate estimation of temperature. The Kalman filter outputs a temperature estimation $\hat{x}_t$ and a standard deviation $\hat{\sigma}_t$, which shows the confidence level (accuracy) of the temperature estimation. These two values are then pipelined to a Federated Learning neural network that outputs 3 signals: fault prediction, optimal charging power and cooling fan activation.
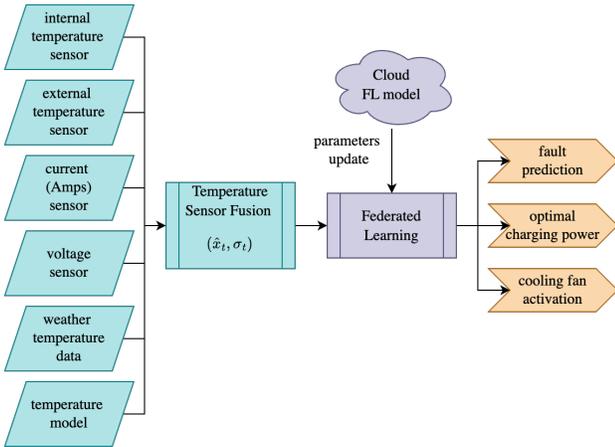

Fig. 1. The proposed system architecture for temperature Sensor Fusion and federated Learning

B. Cloud Level

The Federated Learning framework is organized using Kubernetes nodes and clusters. Each charging point is an autonomous node running a lightweight Kubernetes runtime (k3s) on top of which the sensor fusion and FL algorithms are deployed with Docker containers. Multiple charging points are then grouped together into Kubernetes clusters. In this paper we propose the creation of multiple clusters so that each one of these have a different central Machine Learning model which is unique to the geographic area of its nodes. This means that the charging points that belong to the same cluster, belong to the same geographic area and share common characteristics like weather conditions and temperature and may use the same Energy sub-grid. It is important to use different geographic clusters, because each one will have different probability distributions that the
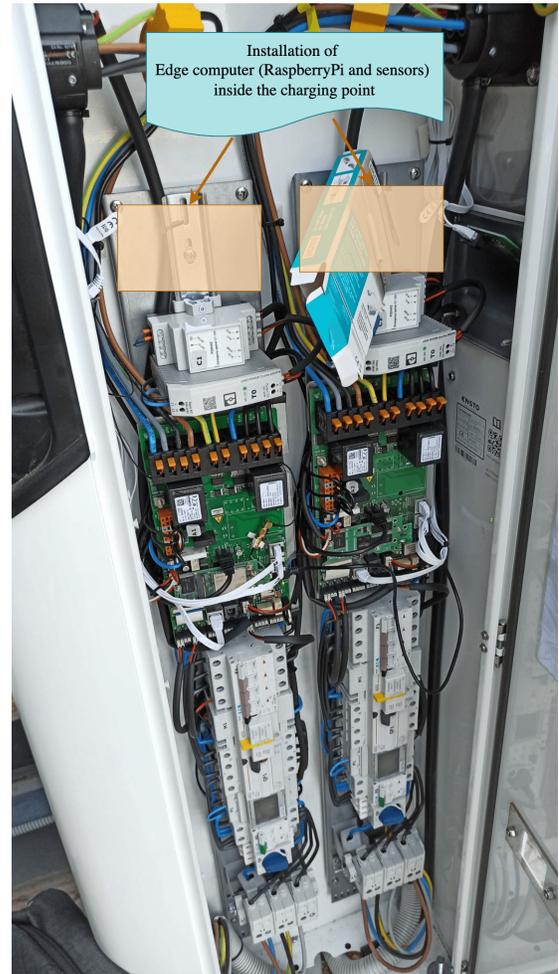

Fig. 2. Installation of the Edge computer and sensors inside the EV Charging Point. In this picture, 2 possible places for mounting the hardware are shown during the installation of the charging point. The charging point shown here is an Ensto Chago Pro EVF200.

Fig. 3. Installation of 5G router in the EV charging point

Federated Learning algorithm is tasked to learn.

At the cloud level, all the locations of each charging point

are stored in a database. Each data point is a pair with latitude and longitude. These locations are then split into clusters. Each cluster of charging point refers to a specific geographical area with similar conditions (weather, same power grid, similar power demand, used mostly by specific people e.g. from a near city, etc.) and has its own Machine Learning model. The clusters are created using the k-means algorithm.

Given a dataset of locations $\mathcal{L} = \{l_1, l_2, \cdots, l_n\}$ where each location $l_i = (lat_i, lon_i)$ represents the latitude and longitude of a point, and a specified number of k geographical clusters, we randomly initialize k cluster centroids $\mathcal{C} = \{c_1, c_2, \cdots, c_k\}$. Assign each data point l_i to the nearest centroid c_j, based on the Euclidean distance:

$$S_j = \{l_i \mid \mathrm{argmin}_j d(l_i, c_j)\} \tag{1}$$

where the Euclidean distance $d(l_i, c_j)$ is given by:

$$d(l_i, c_j) = \sqrt{(lat_i - lat_j)^2 + (lon_i - lon_j)^2} \tag{2}$$

Then on each update step compute the new centroids by averaging the coordinates of all points assigned to each cluster:

$$c_j = \frac{1}{|S_j|} \sum_{l_i \in S_j} l_i \tag{3}$$

We repeat the above steps until the centroids no longer change significantly or a maximum number of iterations is reached. When the clusters are set then each charging point is represented by its centroid and the associated federated learning model.
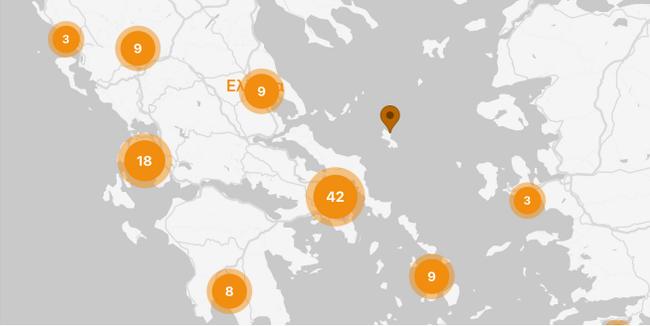

Fig. 4. Edge charging points clustering based on geographical proximity

III. TEMPERATURE SENSOR FUSION

We consider the following state variable that we want to estimate:

$$\mathbf{x}_t = \begin{bmatrix} T_{\text{cable},t} \\ T_{\text{ambient},t} \end{bmatrix} \tag{4}$$

where $T_{\text{cable},t}$ is the temperature of the EV charging cable (at the socket of the EVSE) at time t and $T_{\text{ambient},t}$ is the ambient temperature of the EVSE at time t. We also consider the temperature equations:

$$T_{\text{cable},t} = T_{\text{cable},t-1} + \Delta t \left(k_p P_t - k_c (T_{\text{cable},t-1} - T_{\text{ambient},t-1}) \right) + w_{\text{cable},t} \tag{5}$$

and

$$T_{\text{ambient},t} = T_{\text{ambient},t-1} + w_{\text{ambient},t} \tag{6}$$

where $w_{\text{cable},t} \sim \mathcal{N}(0, Q_{\text{cable}}), w_{\text{ambient},t} \sim \mathcal{N}(0, Q_{\text{ambient}})$ are process noises (Gaussian). We get the following model of temperature states:

$$\mathbf{x}_t = \mathbf{F}\mathbf{x}_{t-1} + \mathbf{B}u_t + \mathbf{w}_t \tag{7}$$

where:

$$\mathbf{F} = \begin{bmatrix} 1 - \Delta t \cdot k_c & \Delta t \cdot k_c \\ 0 & 1 \end{bmatrix}, \quad \mathbf{B} = \begin{bmatrix} \Delta t \cdot k_p \\ 0 \end{bmatrix}, \quad u_t = P_t$$

The calculations for the prediction step of the Kalman filter are the following:

$$\hat{\mathbf{x}}_{t|t-1} = \mathbf{F}_{t-1}\hat{\mathbf{x}}_{t-1|t-1} + \mathbf{B}_{t-1}u_t \tag{8}$$

$$\mathbf{P}_{t|t-1} = \mathbf{F}_{t-1}\mathbf{P}_{t-1|t-1}\mathbf{F}^\top + \mathbf{Q} \tag{9}$$

The calculations for the update step of the Kalman filter are the following:

$$\mathbf{S}_t = \mathbf{H}\mathbf{P}_{t|t-1}\mathbf{H}^\top + \mathbf{R} \tag{10}$$

$$\mathbf{K}_t = \mathbf{P}_{t|t-1}\mathbf{H}^\top \mathbf{S}_t^{-1} \tag{11}$$

$$\hat{\mathbf{x}}_{t|t} = \hat{\mathbf{x}}_{t|t-1} + \mathbf{K}_t \left(\mathbf{y}_t - \mathbf{H}\hat{\mathbf{x}}_{t|t-1} \right) \tag{12}$$

$$\mathbf{P}_{t|t} = (\mathbf{I} - \mathbf{K}_t\mathbf{H})\mathbf{P}_{t|t-1} \tag{13}$$

where $\hat{\mathbf{x}}_{t|t-1} \in \mathbb{R}^{2\times 1}$ is the predicted system state at time step t, before applying the new measurements, $\mathbf{F}_{t-1} \in \mathbb{R}^{2\times 2}$, is the transition matrix that propagates the system state to the next time step , $\mathbf{B}_{t-1} \in \mathbb{R}^{2\times 1}$ is the coefficient vector of the input u_t (control variable, charging power), $\mathbf{P}_{t|t-1} \in \mathbb{R}^{2\times 2}$ is the covariance matrix of the errors of $\mathbf{x}$ at time step t, before applying the new observations and $\mathbf{Q}_{t-1} \in \mathbb{R}^{2\times 2}$ is the process noise, which describes additional uncertainties due to modelling errors.

IV. FEDERATED LEARNING

To leverage the data from many charging stations while preserving privacy, we integrate the above sensor fusion mechanism with Federated Learning (FL). Federated learning is a distributed machine learning approach where each EVSE (charging point) trains a model locally on its own data (e.g. its history of temperature readings and observed faults or adjustments) and periodically exchanges only model parameters or updates with a central server (aggregator). Raw data never leaves the local site, which is beneficial for privacy and bandwidth – especially relevant if chargers are operated by different entities or generate sensitive usage data.

Integration with sensor fusion: In our framework, each charging station is equipped with a lightweight edge computing unit that runs the sensor fusion (Kalman filter) and a machine learning model. After the Kalman filter produces the estimated temperature $\hat{T}_{cable,t}$ and confidence $\hat{\sigma}_t$ at each time step, these values (possibly along with other local features

like current draw or time of day) serve as inputs to the local neural network. This neural network is trained to output three signals critical for optimal charger operation under thermal constraints: (1) Fault prediction, (2) Optimal charging power, and (3) Cooling fan activation.

- **Fault prediction**: a binary or probabilistic output indicating whether an overheating fault is imminent if charging continues under current conditions. This acts as an early warning system. For example, if $\hat{T}_{cable,t}$ is rising rapidly and nearing a safety threshold, the model will predict a fault with high probability before the hardware's built-in cutoff triggers. This prediction allows preemptive action to avoid a hard fault. In essence, the model learns the patterns (temperature trends, ambient conditions) that preceded past thermal shutdowns or failures and can recognize them in advance.
- **Optimal charging power**: a recommended power level (or current limit) to supply to the EV that will prevent overheating while delivering the maximum possible energy. Instead of a binary on/off or a fixed derating, the model can suggest, say, "reduce to 70% power" to keep temperature stable. This output is informed by the estimated temperature and its rate of change. For instance, if the charger is at risk of overheating, the optimal power might be a value P that balances heat generation and dissipation. From the thermal model, one can derive a rough guideline for safe power: requiring that steady-state T_{cable} stays below the max limit T_{max} implies

$$P \approx \frac{k_c}{k_p}\big(T_{\max} - T_{ambient}\big) \qquad (14)$$

(using the thermal equilibrium of the model) – the neural network can learn such relationships and adjust for dynamic conditions.
- **Cooling fan activation**: a binary signal to activate a cooling mechanism (if available) such as a fan when high temperatures are detected. Many high-power chargers include fans or liquid cooling, but even smaller AC chargers can be retrofitted with cooling fans in the enclosure. The model learns to toggle the fan only when needed, i.e. when $\hat{T}_{cable}$ exceeds a certain threshold or is rising quickly. This prevents unnecessary energy use by the fan and ensures cooling is used effectively in tandem with power adjustments. (If a particular charger has no active cooling, this output can be ignored or repurposed for another alert mechanism.)

Each EVSE runs this neural network locally in real time to make immediate control decisions. However, to train this model to be as accurate and generalizable as possible, federated learning is used. Initially, the model can be pre-trained on a generic dataset or simulator. Then, as chargers operate, each one continually improves the model with its own experience: it collects training examples like "[features] -¿ [was there a fault / what power was safe / did we need fan]". Periodically (for example, every hour or at the end of each day or when a significant event occurs), the charger will upload its model updates (e.g. gradients or the new weights) to a cloud server through a 5G connection. The cloud server (central FL aggregator) then performs the Federated Averaging (FedAvg) algorithm to combine updates from all participating chargers:

If there are N chargers in the federation, with charger i holding a local dataset of n_i samples, the global model update can be expressed as:

$$w^{(t+1)} = \sum_{i=1}^{N} \frac{n_i}{\sum_{j=1}^{N} n_j} \, w_i^{(t+1)} \qquad (15)$$

where $w_i^{(t+1)}$ are the local model weights after charger i has performed its local training (at round t), and $w^{(t+1)}$ is the new global model. This weighted averaging means larger local datasets influence the global model more, reflecting their statistical significance. Equivalently, the global model is optimizing a federated objective

$$F(w) = \sum_{i=1}^{N} \frac{n_i}{n} F_i(w) \qquad (16)$$

a weighted sum of each local loss $F_i(w)$. In our scenario, F_i might measure prediction error on charger i's temperature and fault data.

The server then sends the updated global model weights back to each EVSE. Each EVSE updates its local model to this global model and can continue training with new local data, and the cycle repeats. Over many federated rounds, the model learns from the combined experience of all chargers without ever pooling the raw data in one place.

This federated approach is particularly apt for EV infrastructure because data about faults or high-temperature events can be rare and distributed. One charging station in a hot climate might frequently approach thermal limits, while another in a cooler climate rarely does – by sharing a model, the cooler-climate station can still learn how to respond when an extreme heat event occurs from the experiences of the hotter station. Federated learning thus improves the model's generalization across different operating conditions. It has been shown that FL can significantly improve the quality of service of EV charging stations by leveraging data from multiple stations. Our framework extends that notion to improving safety and reliability (not just pricing or service time) via collaborative learning.

V. Conclusion & Future Work

Conclusion: In this paper, we presented a novel framework for optimal EV charging point operation and fault prevention by combining 5G-enabled sensor fusion with federated learning. The approach targets one of the critical practical issues in EV infrastructure – overheating of charging equipment – and provides a solution that is both local (real-time control at each charger) and global (collaborative learning across the network). By deploying a Kalman filter-based sensor fusion algorithm at the edge, we obtain accurate real-time estimates of the charger temperature and its rate of change. This feeds

into a federated learning model – a neural network distributed across all charging stations – which predicts impending faults, recommends safe charging power levels, and controls cooling mechanisms. The integration of 5G connectivity ensures that model updates and alerts are communicated instantly, enabling a true IoT system of smart chargers.

The research validated the technical soundness of each component: the Kalman filter improved temperature sensing fidelity, and the federated neural network, trained on data from multiple chargers, outperformed isolated models in predicting and preventing faults. Case study simulations demonstrated that our system can successfully prevent EVSE thermal faults, avoiding charging session interruptions by dynamically adjusting power output and cooling. This leads to a more reliable charging service, particularly in extreme ambient conditions where conventional chargers might fail. The use of federated learning means these benefits are achieved without centralizing sensitive operational data, and each charging point continuously becomes "smarter" as more data (from itself and peers) becomes available. In essence, the work bridges the gap between robust control (via sensor fusion and real-time actuation) and distributed AI (via FL) in the context of EV charging. It also showcases how 5G networks can empower edge devices (like EVSEs) to participate in advanced algorithms that were previously confined to centralized systems, by providing the needed bandwidth and latency guarantees for such distributed learning.

Improvements over current state of the art: Traditional EVSE safety mechanisms are reactive – they monitor temperature and shut down or derate once a threshold is exceeded. Our framework adds intelligence and proactivity: it predicts when a threshold would be exceeded and acts to prevent it, and it continually learns optimal responses from a fleet of chargers. This approach improves on earlier works by incorporating both model-based estimation (Kalman filtering) and data-driven learning in a hierarchical manner. Compared to a purely centralized solution, our distributed method is more scalable and respects user privacy. Compared to standalone smart chargers, our collaborative approach achieves higher accuracy and adaptability. The successful deployment in a Kubernetes cluster spanning multiple edge devices also underscores the practical feasibility of the system in real-world settings (e.g., a network operator could roll this out via over-the-air updates to chargers).

Future Work: This study opens several avenues for further research and development:

Advanced Machine Learning Models: The current neural network could be enhanced by more sophisticated architectures. A natural extension is to use a recurrent neural network (RNN) or LSTM that observes a sequence of past temperature estimates and perhaps power levels to better predict future faults (capturing temporal patterns explicitly). Another idea is to treat the power adjustment as a reinforcement learning problem: the system could learn an optimal policy (with reward for keeping the charger active but penalties for high temperature) via federated reinforcement learning. Recent work in

EV charging control with multi-agent reinforcement learning shows potential, and combining that with FL could allow stations to learn control policies cooperatively. We can also consider multi-task learning where the fault prediction, power recommendation, and cooling control are separate outputs – tuning the loss function weighting for these tasks might improve, for example, the fault prediction accuracy without sacrificing power control quality.

Wider Range of Operating Conditions: So far, our results are strong for high-temperature scenarios. In the future, we plan to ensure the system also gracefully handles other stress conditions. For example, very cold weather might not be a thermal risk, but could impact battery acceptance and thus indirectly the charger's behavior (though primarily a battery issue). Another case is electrical faults or grid fluctuations – while our current model focuses on thermal faults, a similar approach could be used for other fault types (over-voltage, ground faults, etc.), by fusing relevant sensor data and training models to predict those events. This would broaden the scope from purely thermal management to a more comprehensive EVSE health monitoring system.

Federated Learning Enhancements: Deploying FL in a large, real-world EV charging network will pose challenges that we aim to address. One challenge is imbalanced and non-IID data – some chargers will contribute many updates (busy stations) while others are seldom used. We plan to experiment with advanced FL aggregation schemes (such as FedAvgM, FedProx, or FedNova) and weighting strategies that account for data quality, not just quantity. Techniques like the FedAvg-ContribData algorithm could be applied to weight updates by their usefulness, which might help in our case where "rare events" (overheating in cool regions) are underrepresented. We will also look into personalized federated learning: after the global model is learned, each charger can have a fine-tuning step on its own data to slightly personalize the model to its environment (while still benefiting from the global knowledge). This could yield a slight performance boost, especially for outlier locations.

Real-World Pilot and User Study: A crucial next step is a pilot deployment on actual charging stations in the field. This would validate assumptions (e.g. communication reliability over 5G in various locations, integration with existing charger firmware) and measure real-world impact. We also need to study the user acceptance: how do EV owners perceive dynamic charging rate adjustments? It will be important to communicate that a slowdown is preventive and actually avoids a worse outcome. Perhaps the system could interface with the EV or mobile app to notify the user "Charging power reduced to prevent overheating, to avoid an interruption." Such transparency can improve user trust. In the future, EVs might even share their battery temperature or state with the charger (via protocols like ISO 15118), allowing a cooperative thermal management between vehicle and charger – for example, an EV could decide to draw a lower current if it knows the station is overheating or vice versa. Exploring this vehicle-to-charger communication for thermal management is an exciting

direction.

Scalability and Edge Analytics: As the number of chargers grows, the federated learning infrastructure must scale. Network slicing in 5G could be used to dedicate bandwidth to critical FL updates, ensuring timely model convergence even under heavy network traffic. We will also investigate if hierarchical FL can reduce communication: e.g., first aggregate models at a regional level (edge cloud) then a global level, to minimize latency. Additionally, incorporating anomaly detection in the model update process is important for security – if one charger sends anomalous updates (due to a malfunction or attack), the aggregator should detect and ignore it (using techniques like update clustering or blockchain verification).

REFERENCES

[1] S. Hamdare, O. Kaiwartya, M. Jugran, D. J. Brown, and P. Vyas, "Analysis of ev charging infrastructure and its impact on public adoption: Examining the critical role of charging stations in the acceleration of electric vehicle market growth," *Proceedings of the 16th International Conference on PErvasive Technologies Related to Assistive Environments*, 2023. [Online]. Available: https://api.semanticscholar.org/CorpusID:260776983

[2] Y. Zhang, J. Chen, L. X. Cai, and J. Pan, "Expanding ev charging networks considering transportation pattern and power supply limit," *IEEE Transactions on Smart Grid*, vol. 10, pp. 6332–6342, 2019. [Online]. Available: https://api.semanticscholar.org/CorpusID:115877469

[3] N. Al-Dahabreh, M. A. Sayed, K. Sarieddine, M. K. Elhattab, M. J. Khabbaz, R. Atallah, and C. M. Assi, "A data-driven framework for improving public ev charging infrastructure: Modeling and forecasting," *IEEE Transactions on Intelligent Transportation Systems*, vol. 25, pp. 5935–5948, 2023. [Online]. Available: https://api.semanticscholar.org/CorpusID:266163098

[4] B. Borlaug, Y. Ge, D.-Y. Lee, Z. Liu, E. Rose, L. S. Luhring, M. S. Louis-Sanchez, E. Wood, and F. Yang, "Modeling future demand for ev charging infrastructure," 2023. [Online]. Available: https://api.semanticscholar.org/CorpusID:266247649

[5] J. Shepard, "What are common sensor fusion applications in EVs? — evengineeringonline.com," https://www.evengineeringonline.com/what-are-common-sensor-fusion-applications-in-electric-and-autonomous-vehicles/, [Accessed 16-11-2024].

[6] M. Yasin, "Enhancing ev charging networks: Advanced fusion techniques with insights from lstm, bayesian networks, and deep learning," *International Journal of Intelligent Systems and Applications in Engineering*, vol. 12, no. 3, p. 3363–3370, Mar. 2024. [Online]. Available: https://ijisae.org/index.php/IJISAE/article/view/5972

[7] R. Lei, Y. B. Guo, J. Yan, and W. uo, "Bridging data gaps: A federated learning approach to heat emission prediction in laser powder bed fusion," *Journal of Manufacturing Science and Engineering*, vol. 146, no. 10, p. 101002, 07 2024. [Online]. Available: https://doi.org/10.1115/1.4065888