https://doi.org/10.2352/EI.2026.38.3.MOBMU-327

Deep learning based vehicle classification: Detecting EVs and gasoline cars in Berlin using convolutional neural networks

Raghav Tandon [1], *Hamid Mostofi* [1], *Navaneeth Shivananjappa* [1], *Reiner Creutzburg* [1,2,3]

[1] *SRH University of Applied Sciences Heidelberg, Campus Berlin, School of Technology and Architecture, Sonnenallee 221c, D-12059 Berlin, Germany*
[2] *Technische Hochschule Brandenburg, Department of Informatics and Media, Magdeburger Str. 50, D-14770 Brandenburg, Germany*
[3] *German University of Applied Science, Marlene-Dietrich-Allee 14, D-14482 Potsdam, Germany*

Email: raghavtandonsrh@gmail.com, hamid.mostofi@srh.de, navaneeth.shivananjappa@srh.de, reiner.creutzburg@srh.de, creutzburg@th-brandenburg.de, reiner.creutzburg@german-uds.de

Keywords: Deep Learning, Vehicle Classification, Electric Vehicles (EVs), Convolutional Neural Networks, Object Detection, Smart Parking, Edge Computing, Privacy-Aware Systems

Abstract

The rapid growth of electric vehicles (EVs) has introduced new challenges for urban parking management, mainly in enforcing EV-designated parking spaces without intrusive infrastructure. This paper presents a deep-learning-based vision system for the automated classification of electric and gasoline vehicles in urban parking environments, using convolutional neural networks trained on real-world data from Berlin, Germany. A YOLO-based object detection model is employed to identify visually distinctive EV-specific features in rear-view vehicle images while preserving privacy by anonymizing license plates. The proposed approach relies solely on visual cues, eliminating the need for vehicle metadata, sensors, or network connectivity. Experimental results demonstrate robust classification performance, achieving high detection accuracy and consistent results across desktop and edge computing platforms. To validate real-world applicability, the trained model is deployed on both a mobile device and a low-cost Raspberry Pi-based edge system, enabling fully offline operation. These results indicate that deep learning-based visual classification can provide a scalable, privacy-aware solution for smart parking systems and urban mobility applications. This supports the effective management of EV infrastructure in modern cities.

Introduction

Rapid adoption of EVs has transformed urban mobility worldwide and is driven by environmental regulations, sustainability goals, and advances in vehicle electrification. Expanding EV infrastructure in cities, such as charging stations and EV-designated parking spaces, has introduced practical challenges, including enforcing proper use of these facilities. Conventional enforcement methods typically rely on manual inspection, vehicle registration databases, or sensor-based solutions, which are either labor-intensive, intrusive, or costly to deploy at scale. Consequently, there is a growing need for automated, low-cost, and privacy-aware systems capable of identifying EVs in real-world urban environments.

Recent advances in computer vision and deep learning enable robust vehicle detection and classification from visual data alone. CNNs are robust to diverse lighting, weather, and occlusion conditions, making vision-based vehicle classification a promising alternative to infrastructure-heavy or metadata-dependent methods. Distinguishing electric from gasoline vehicles remains difficult because their exteriors often appear similar, and urban deployment must respect privacy, computational limits, and offline operation.

This paper proposes a deep-learning-based vehicle classification framework that distinguishes between electric and gasoline vehicles using visual cues in rear-view images captured in urban parking areas. The approach is trained and evaluated on a Berlin-specific dataset that reflects realistic urban conditions and vehicle distributions. A YOLO-based detector identifies EV-specific visual indicators while preserving privacy through automatic license plate anonymization. Unlike sensor- or database-driven systems, the method relies solely on image data and operates without network connectivity or access to vehicle registration data.

Beyond model development and evaluation, this work focuses on real-world deployment on resource-constrained platforms. The trained model is validated on a mobile device and a low-cost Raspberry Pi-based edge system, demonstrating feasible offline inference for smart parking and urban mobility. The paper's key contributions are: (1) a vision-based deep learning framework for EV and gasoline vehicle classification in urban parking, (2) an experimental evaluation using real-world data from Berlin, and (3) validation through deployment on mobile and edge platforms, underscoring scalability, privacy awareness, and suitability for smart city infrastructure.

Literature Review

Vision-based vehicle detection and classification have been extensively studied in the context of intelligent transportation systems, traffic monitoring, and smart city applications. Early approaches relied on handcrafted features combined with traditional machine learning classifiers; however, these methods often struggled under varying illumination, occlusions, and complex urban backgrounds. With advances in deep learning, CNNs have become the dominant approach for vehicle detection and recognition tasks due to their ability to learn hierarchical visual features from large-scale datasets automatically [1] [2].

Recent studies have demonstrated the effectiveness of CNN-based object detection frameworks, including region-based and single-stage detectors, for accurately localizing and classifying vehicles in real-world traffic scenes [3]. Among these, single-stage detectors have gained attention for their real-time performance and suitability for deployment in constrained environments. These models have been successfully applied to tasks including vehicle counting, type classification, and traffic flow analysis in urban settings [4].

The specific problem of distinguishing electric vehicles from gasoline-powered vehicles using visual data remains relatively underexplored. Several existing systems identify EVs indirectly, using license plate databases, onboard sensors, radio-frequency identification (RFID), or integration with charging infrastructure [5]. While effective, these approaches often require access to external data sources, network connectivity, or specialized hardware, raising concerns related to scalability, cost, and user privacy. Vision-based identification using external visual cues offers a non-intrusive alternative but introduces challenges due to the visual similarity between EVs and conventional vehicles.

More recent research has explored the use of deep learning models to detect EV-specific visual indicators, such as badges, logos, or design elements, using image-based classification techniques [6], [7]. These approaches demonstrate promising accuracy but are typically evaluated under controlled conditions or rely on high-resolution imagery. Furthermore, privacy considerations, particularly in public parking environments, are often insufficiently addressed, despite increasing regulatory emphasis on data protection in urban surveillance systems [8].

In parallel, there has been growing interest in deploying deep learning models on edge devices to reduce latency, reduce dependence on cloud services, and mitigate exposure of sensitive data. Studies have shown that object detection models can be adapted for execution on resource-constrained platforms such as single-board computers and mobile devices, albeit with tradeoffs in inference speed and model complexity [9]. Such edge-based deployments are particularly relevant for smart parking applications, where offline operation and low-cost infrastructure are critical.

Building on prior work, this paper presents a vision-based deep learning approach for classifying electric and gasoline vehicles in urban parking environments using real-world data. Unlike database-driven or sensor-based systems, the proposed method relies solely on visual information, explicitly incorporates privacy-preserving mechanisms, and validates feasibility through mobile and edge-device deployments.

Methodology

This section describes the design and implementation of the proposed deep-learning-based vehicle classification framework, including dataset preparation, model architecture, training strategy, and privacy-preserving measures for urban deployment.

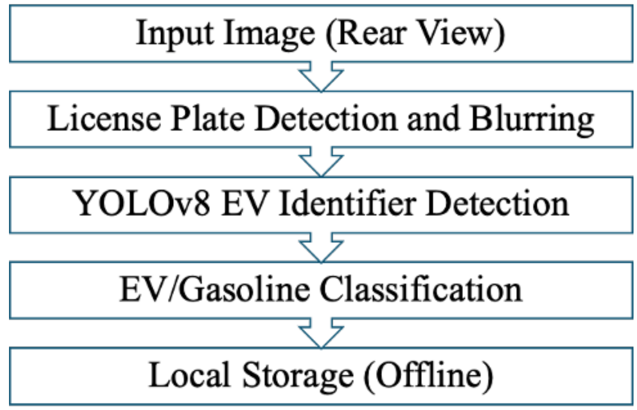

Figure 1. *Overview of the proposed deep learning–based vehicle classification framework, illustrating privacy-preserving preprocessing and offline inference*

Dataset Development and Preprocessing

The dataset used in this study was developed specifically for urban parking environments in Berlin, Germany, and captures real-world variations in vehicle types, lighting conditions, viewing angles, and occlusions. Images were collected primarily from rear-facing perspectives of parked vehicles, as this view consistently exposes visual identifiers relevant for differentiating electric vehicles from gasoline-powered vehicles. The dataset includes both electric and gasoline vehicles, with careful attention to maintaining representative coverage of vehicle designs commonly observed in Berlin.

All images were manually annotated using bounding boxes to label EV-specific visual identifiers. To improve robustness and generalization, standard preprocessing and data augmentation techniques were applied, including image resizing, normalization, and controlled geometric and photometric transformations. These steps ensured that the model could handle diverse urban conditions, including partial occlusions, varying illumination, and non-uniform parking orientations.

Model Architecture

A YOLOv8-based convolutional neural network was selected for this study due to its balance between detection accuracy and computational efficiency. The model follows a single-stage object detection paradigm, enabling end-to-end learning of both localization and classification tasks within a unified framework. This design choice is particularly suitable for real-time or near-real-time inference in constrained environments such as mobile devices and edge computing platforms.

The network architecture comprises a feature-extraction backbone, a feature-aggregation neck, and a detection head optimized to identify EV-specific visual cues. Rather than classifying entire vehicles based on global appearance, the model focuses on localized visual indicators associated with electric vehicles, improving robustness against stylistic similarities between EVs and gasoline vehicles.

Training Strategy

Model training was performed using supervised learning on the annotated dataset, with the data split into training, validation, and test subsets. The training process employed transfer learning, initializing the network with pretrained weights to accelerate convergence and improve performance on limited domain-specific data. Hyperparameters such as learning rate, batch size, and number of training epochs were selected based on empirical validation performance.

Evaluation metrics included mean Average Precision (mAP), precision, recall, and F1-score, with particular emphasis on mAP at an intersection-over-union threshold of 0.5. These metrics provide a comprehensive assessment of both detection accuracy and localization quality. Validation performance was monitored throughout training to mitigate overfitting and ensure model generalization.

Privacy Preserving Measures

Given the deployment context in public urban spaces, privacy protection was treated as a core design requirement. All collected images were processed to anonymize personally identifiable information, particularly vehicle license plates. A dedicated license plate detection and optical character recognition pipeline was applied to identify and blur license plate regions before storage or further processing, ensuring compliance with data protection regulations.

Importantly, all inference operations were designed to run locally without reliance on cloud services or external databases. This offline processing strategy minimizes data exposure, reduces network dependency, and enhances system trustworthiness for real-world deployment in smart city environments.

Deployment Readiness

The methodology was designed with deployment feasibility as a central consideration, ensuring that the proposed framework extends beyond theoretical evaluation into practical, real-world applicability. The trained model was validated not only on desktop hardware, where baseline performance could be established, but was also carefully adapted for execution on mobile devices and resource-constrained edge-computing platforms. This cross-platform validation demonstrates the approach's flexibility and its ability to operate under varying computational constraints.

Particular emphasis was placed on maintaining consistent classification performance as the system transitioned from high-performance desktop environments to low-cost, decentralized hardware. This ensures that the system can be deployed in realistic urban settings without requiring specialized infrastructure or continuous cloud connectivity. As a result, the proposed approach remains scalable, cost-effective, and well-suited for decentralized smart parking systems, where low latency, offline operation, and data privacy are essential requirements.

Experiments and Results

This section presents the experimental setup and quantitative evaluation of the proposed deep-learning-based vehicle classification framework. The objective of the experiments was to assess the model's ability to reliably distinguish electric vehicles from gasoline-powered vehicles in real-world urban parking environments.

Experimental Setup

The YOLOv8-based model was trained and evaluated using a dataset split into training (80%), validation (10%), and test (10%) subsets. All experiments were conducted under fully offline conditions to align with the intended deployment in privacy-sensitive urban environments. Initial evaluations were performed on a desktop computing system to establish baseline performance, followed by inference on resource-constrained platforms to assess deployment feasibility.

All evaluations were performed using rear-view images of parked vehicles, reflecting realistic parking scenarios where vehicle motion is absent and visual identifiers may be partially occluded or inconsistently visible. Identical preprocessing and inference pipelines were maintained across platforms to ensure fair comparison of results.

Evaluation Metrics

Model performance was evaluated using standard object detection and classification metrics, including precision, recall, F1-score, and mean Average Precision (mAP) at Intersection over Union at 0.5 and 0.5-0.95. Precision measures the proportion of correctly identified positive detections, while recall reflects the model's ability to identify all relevant instances within the dataset. The F1-score provides a harmonic balance between precision and recall, offering a single metric to assess overall classification performance. These metrics collectively provide a comprehensive evaluation of detection reliability and classification robustness across varying urban conditions and image complexities, thereby improving performance.

Mean Average Precision at an intersection-over-union threshold of 0.5 (mAP@0.5) was used as the primary evaluation metric, as it provides a balanced measure of detection accuracy and localization quality for urban applications. This metric evaluates both the correctness of predicted bounding boxes and their alignment with ground truth annotations, making it particularly suitable for real-world scenarios where precise localization of EV-specific visual features is required.

Quantitative Results

The trained model achieved a mAP@0.5 of approximately 88.9% on the test dataset, demonstrating strong classification performance in distinguishing electric vehicles from gasoline-powered vehicles under real-world urban conditions. Comparable performance was observed across the training and validation subsets, indicating effective generalization and limited overfitting.

Precision and recall remained consistently high, confirming that the model detected EV-specific visual identifiers with low false-positive and false-negative rates. The model performed robustly across diverse lighting conditions, partial occlusions, and non-uniform parking orientations commonly encountered in urban environments.

Misclassifications were primarily observed when EV-specific visual cues, such as charging ports or distinctive badging, were either fully occluded or visually ambiguous. These cases highlight inherent limitations of purely image-based classification when visual identifiers are not clearly observable.

Figure 2. Example detection results showing EV-specific visual identifiers detected from rear-view vehicle images under urban conditions

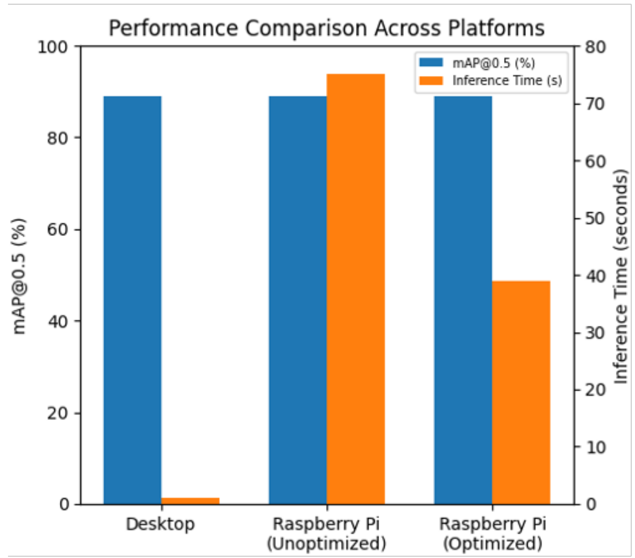

Figure 3. Comparison of classification accuracy (mAP@0.5) and inference time across desktop and edge-computing platforms

Platform Performance Comparison

To evaluate the feasibility of real-world deployment, the trained model was executed on both desktop hardware and a low-cost edge computing platform based on a Raspberry Pi 4B. Desktop inference achieved rapid processing suitable for near-real-time analysis, serving as a performance baseline.

On the edge platform, un-optimized inference times ranged between 70 and 80 seconds per image. After applying hardware optimizations, including CPU overclocking and active cooling, the inference time was reduced to approximately 38 – 40 seconds per image, while maintaining classification accuracy comparable to desktop inference (88.9% mAP@0.5). This demonstrates that the proposed model can operate reliably in constrained environments without significant loss of detection performance.

In addition to quantitative evaluation on desktop and edge platforms, the trained model was deployed and validated in a mobile application environment, demonstrating portability and practical usability beyond controlled experimental settings.

Discussion of Results

The experimental results provide quantitative evidence that deep-learning-based visual classification is a viable and effective solution for identifying electric vehicles in urban parking environments. The achieved classification accuracy, combined with successful offline execution on low-cost hardware, demonstrates the practicality of the proposed approach for decentralized smart parking systems.

By eliminating dependence on external vehicle databases, cloud infrastructure, or intrusive sensors, the system provides a scalable, privacy-aware alternative to traditional enforcement mechanisms. The results highlight the potential of CNN-based approaches for smart city applications, where accuracy, privacy compliance, and deployment flexibility are critical requirements.

Prototype Deployment and Discussion

To validate the practical applicability of the proposed deep learning based vehicle classification framework beyond controlled experimental evaluation, the trained model was deployed on two real-world prototype systems: a mobile application and a low-cost edge computing platform. These deployments were not intended to redefine model accuracy metrics but to demonstrate feasibility, portability, and privacy-aware operation in realistic urban settings.

Mobile Application Prototype

A mobile application prototype was developed to demonstrate on-device vehicle classification using the trained convolutional neural network. The application was implemented using the Flutter framework with Dart for cross-platform development and deployed on an iPhone 12 mini for testing. The primary objective of this prototype was to evaluate whether the trained YOLOv8-based model could be integrated into a mobile workflow and perform inference using images captured directly from a smartphone camera.

The application enables users to capture still images of parked vehicles from the rear-view perspective. Once an image is captured, it is processed through the trained model to detect EV-specific visual identifiers. Based on the inference result, the image is classified as either an electric or gasoline vehicle and stored locally in the corresponding folders. All detected images include bounding boxes and confidence scores for identified EV indicators. The interface is designed to be intuitive and lightweight, allowing users to quickly capture, process, and review classification results without requiring technical expertise or additional configuration.

Privacy considerations were explicitly incorporated into the

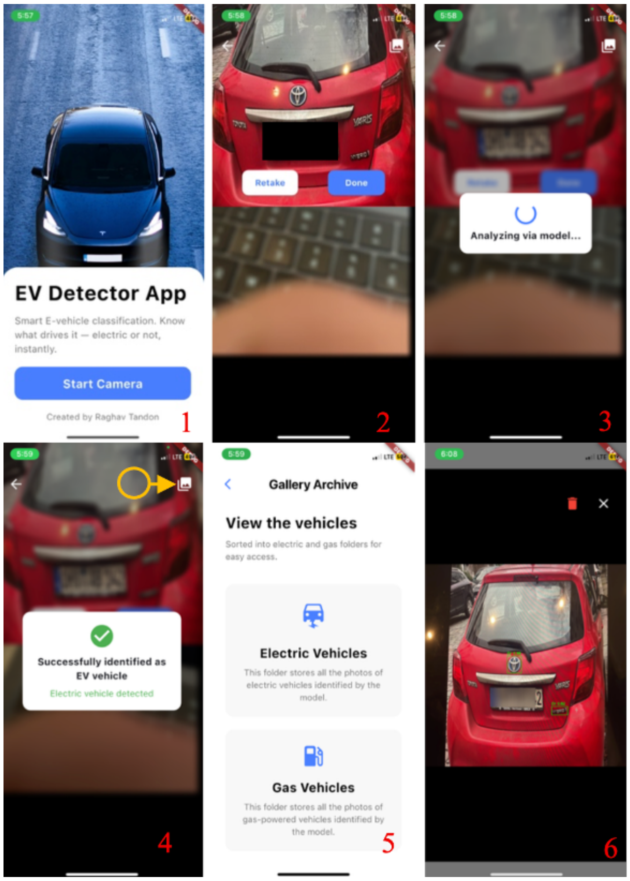

Figure 4. *Mobile application prototype demonstrating on-device EV and gasoline vehicle classification using the trained CNN model*

mobile prototype. License plates were automatically anonymized before storage, ensuring that no personally identifiable information was retained. All inference operations were conducted locally without reliance on cloud services or remote servers, demonstrating the feasibility of privacy-preserving, offline mobile deployment. This design approach minimizes data exposure risks and aligns with regulatory requirements for handling visual data in public urban environments.

While the mobile application successfully demonstrated image-based classification functionality, certain limitations remain. The prototype currently supports inference on still images only and does not perform real-time video-based detection. Additionally, full native on-device inference using platform-specific model formats has not yet been implemented, requiring further work on model conversion and optimization for mobile deployment. Future improvements may include integration with platform-specific acceleration frameworks to enhance inference speed and enable real-time processing capabilities.

Raspberry Pi–Based Edge Computing Prototype

In addition to the mobile application, the trained model was deployed on a standalone edge computing system using a Raspberry Pi 4B with 4 GB RAM. This prototype was designed to simulate a smart parking sensor that identifies the electric vehicles in an offline environment without network connectivity or cloud infrastructure.

The system utilizes a camera module to capture images of parked vehicles, which are then processed locally using a two-stage inference pipeline. In the first stage, license plates are detected and anonymized using optical character recognition and image processing techniques. In the second stage, the anonymized image is passed to the trained YOLOv8-based model to detect EV-specific visual identifiers. Based on the detection outcome, images are categorized as electric or gasoline vehicles and stored locally.

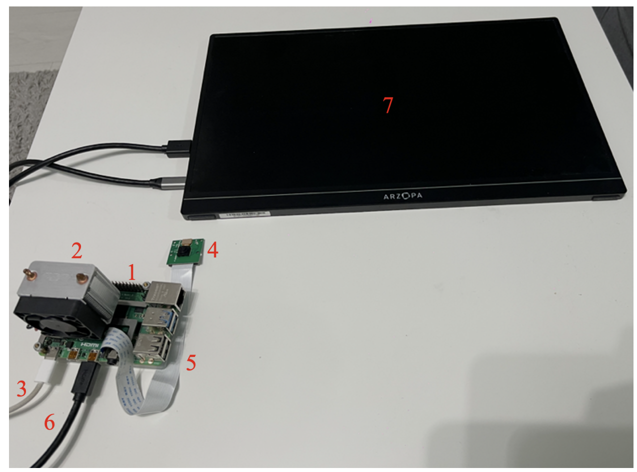

Figure 5. *Raspberry Pi Edge Computing Prototype*

Figure 5 explains the setup for the edge computing prototype. An overclocked Raspberry Pi (1) with active cooling (2) was used to power the prototype through a 15W cable (3). The Raspberry Pi camera takes the photos (4) connected through a ribbon cable (5). The output was transmitted via a micro-HDMI cable (6) to the external screen (7).

Initial inference times on the edge platform ranged between 70 and 80 seconds per image. After applying hardware optimizations, including CPU overclocking and active cooling, inference times were reduced to approximately 38–40 seconds per image while maintaining classification accuracy comparable to that of desktop evaluation. Although inference latency remains higher than that of desktop systems, the achieved performance is sufficient for smart parking applications where immediate real-time response is not mandatory.

The edge prototype operates entirely offline, ensuring data privacy and eliminating exposure to network-based security risks. This can be extended to interface with external hardware components, such as access barriers or indicator systems, enabling automated enforcement of EV-designated parking spaces.

Practical Implications and Limitations

The successful deployment of the trained model on both mobile and edge platforms demonstrates the portability and robustness of the proposed approach. These prototypes validate that deep-learning-based visual classification of electric vehicles can be performed outside laboratory environments using low-cost, decentralized hardware.

However, several limitations must be acknowledged. Inference latency on resource-constrained devices remains a challenge and may require further optimization through model compression, quantization, or pruning. Additionally, image-based classification relies on the visibility of EV-specific visual cues, which may be obscured under certain parking conditions. Despite these limitations, the prototypes confirm the feasibility of integrating the proposed framework into real-world smart parking and urban mobility systems.

Future Work

This paper presents a deep learning framework for classifying electric and gasoline vehicles in urban parking using convolutional neural networks and only rear-view images. It overcomes limitations of traditional methods that require external databases, specialized sensors, or cloud infrastructure. Evaluated on real-world data from Berlin, it achieves a mean Average Precision of about 88.9% under diverse urban conditions.

The approach was also validated through deployment on mobile and edge platforms. Integrating the trained model into a mobile app and a low-cost Raspberry Pi system demonstrates its portability, scalability, and privacy awareness. Offline operation and automatic license plate anonymization before saving any information support data protection compliance, making this framework suitable for real-world smart parking and urban mobility applications.

Future work shall improve inference efficiency via model optimization (e.g., quantization and pruning) and enable real-time video-based detection. Advancements in native on-device inference for mobile platforms can enhance the usability and performance of this model. Integrating contextual data sources, such as parking infrastructure or other sensors, could strengthen classification robustness. Expanding the dataset to cover more vehicle models and urban environments will improve generalization and support large-scale deployment. This proposed framework offers a foundation for privacy-aware, vision-based vehicle classification systems that support intelligent transportation and sustainable urban mobility.

References

[1] R. Abiyev and M. Arslan, "Vehicle detection systems for intelligent driving using deep convolutional neural networks," Discov. Artif. Intell., vol. 3, no. 1, p. 16, May 2023, doi: 10.1007/s44163-023-00062-8.

[2] [2] N. M. Alahdal, F. Abukhodair, L. H. Meftah, and A. Cherif, "Real-time Object Detection in Autonomous Vehicles with YOLO," Procedia Comput. Sci., vol. 246, pp. 2792–2801, 2024, doi: 10.1016/j.procs.2024.09.392.

[3] [3] M. K. Alam et al., "Faster RCNN based robust vehicle detection algorithm for identifying and classifying vehicles," J. Real-Time Image Process., vol. 20, no. 5, p. 93, Oct. 2023, doi: 10.1007/s11554-023-01344-1.

[4] [4] M. A. Berwo et al., "Deep Learning Techniques for Vehicle Detection and Classification from Images/Videos: A Survey," Sensors, vol. 23, no. 10, p. 4832, May 2023, doi: 10.3390/s23104832.

[5] [5] V. Agarwal and G. Bansal, "Automatic number plate detection and recognition using YOLO world," Comput. Electr. Eng., vol. 120, p. 109646, Dec. 2024, doi: 10.1016/j.compeleceng.2024.109646.

[6] [6] J. Zhang, S. Yang, C. Bo, and Z. Zhang, "Vehicle logo detection based on deep convolutional networks," Comput. Electr. Eng., vol. 90, p. 107004, Mar. 2021, doi: 10.1016/j.compeleceng.2021.107004.

[7] [7] F. Zhang, C. Li, and F. Yang, "Vehicle Detection in Urban Traffic Surveillance Images Based on Convolutional Neural Networks with Feature Concatenation," Sensors, vol. 19, no. 3, p. 594, Jan. 2019, doi: 10.3390/s19030594.

[8] [8] S. Bohn and J.-C. Rogge, "The framing of green innovations—a comparative topic modeling study on the public frames of the electric vehicle in Germany and UK," J. Clean. Prod., vol. 364, p. 132499, Sep. 2022, doi: 10.1016/j.jclepro.2022.132499.

[9] [9] Z. Charouh, A. Ezzouhri, M. Ghogho, and Z. Guennoun, "A Resource-Efficient CNN-Based Method for Moving Vehicle Detection," Sensors, vol. 22, no. 3, p. 1193, Feb. 2022, doi: 10.3390/s22031193.

Author Biography

Raghav Tandon is an automotive and mobility engineer whose work lies at the intersection of intelligent transportation systems, computer vision, and sustainable urban mobility. His research focuses on applying deep learning and data-driven methods to real-world mobility challenges, with particular emphasis on privacy-aware vehicle classification and smart parking applications. He has experience translating research concepts into deployable mobile and edge-based systems for urban environments.

Hamid Mostofi Darbani is a Senior Researcher at the Faculty of Electrical Engineering and Computer Science, TU Berlin. His research focuses on applying data science and artificial intelligence to sustainability and smart mobility, with particular emphasis on social acceptance, perception, and socio-economic impacts. He received his PhD (summa cum laude) from TU Berlin in 2021. He is also a part-time Professor of Data Science and Artificial Intelligence at SRH University of Applied Sciences, Heidelberg (Berlin Campus).

Navaneeth Shivananjappa is a Lecturer at SRH University of Applied Sciences, Berlin School of Technology and Architecture, specializing in Web Application Penetration Testing and Cybersecurity. His research interests include Cybersecurity, Cybersecurity tools, Penetration Testing, Web Application Security, Kubernetes Security, Cloud Security, and Cybersecurity Awareness.

Reiner Creutzburg is a Retired Professor of Applied Computer Science at the Technische Hochschule Brandenburg in Brandenburg, Germany. Since 2019, he has been a Professor of IT Security at the SRH Berlin University of Applied Sciences. In 2025, he was appointed as a Senior Professor of Cybersecurity at the newly founded German University of Digital Science in Potsdam, Germany. He is a member of the IEEE and SPIE, and has served as chairman of the Multimedia on Mobile Devices (MOBMU) Conference at the Electronic Imaging conferences since 2005. In 2019, he was elected a member of the Leibniz Society of Sciences to Berlin e.V. His research interests include Cybersecurity, Digital Forensics, Open-Source Intelligence (OSINT), Multimedia Signal Processing, e-learning, and Modern Digital Media and Imaging Applications.