

AI-Enabled Accessible Parking Monitoring and Enforcement: A Comprehensive Review of Computer Vision, ALPR, IoT, and Edge Intelligence

Navdeep Kaur¹, Deepti Sharda¹, Ritika Bansal¹, Deeksha Gupta^{1*}, Mandeep K. Chawla¹

¹ Mehr Chand Mahajan DAV College for Women, Chandigarh, India

*Corresponding author: deeksha.784@gmail.com

Abstract: Accessible parking is a critical element of inclusive mobility because designated spaces are intended to reduce physical and logistical barriers for people with disabilities. However, unauthorized occupation, misuse of permits, inadequate monitoring, and intermittent manual enforcement can reduce the practical availability of these spaces. Recent progress in artificial intelligence (AI), computer vision, automatic license plate recognition (ALPR), optical character recognition (OCR), Internet of Things (IoT), edge computing, and cloud platforms creates an opportunity to transform accessible-parking management from periodic inspection into continuous, evidence-assisted monitoring. This comprehensive review examines the technological components required for intelligent accessible-parking monitoring, including accessible-space identification, vehicle detection, parking occupancy analysis, license-plate recognition, accessibility credential detection, multi-object tracking, authorization verification, temporal reasoning, IoT communication, edge/cloud deployment, privacy, security, and human supervision. Unlike conventional smart-parking solutions that primarily determine whether a space is occupied, accessible-parking enforcement requires the integration of spatial, visual, temporal, identity, and authorization evidence. The reviewed literature indicates that individual components such as parking occupancy detection, object detection, ALPR, OCR, tracking, and IoT architectures have reached substantial technical maturity, whereas complete accessible-parking enforcement remains comparatively less developed. Existing datasets are also fragmented: general parking datasets provide strong support for occupancy research but generally lack accessibility-credential and authorization annotations. Directly relevant work demonstrates the feasibility of combining vehicle, plate, and disability-badge recognition with authorization services, while recent aerial-imagery research demonstrates the possibility of mapping accessible parking at scale. Based on these observations, this review presents an evidence-centric architecture that connects accessible-space mapping, vehicle detection, tracking, plate recognition, credential analysis, authorization databases, temporal reasoning, confidence estimation, secure event storage, and human review. The review further identifies challenges involving small and occluded credentials, environmental variation, cross-site generalization, privacy, legal variation, uncertainty, and the absence of standardized end-to-end benchmarks. The paper concludes with research directions focused on multimodal datasets, privacy-preserving edge intelligence, uncertainty-aware decision making, explainability, cross-condition validation, and human-centered enforcement.

Keywords: accessible parking; smart parking; artificial intelligence; computer vision; automatic license plate recognition; OCR; IoT; edge AI; parking enforcement; accessibility; smart cities.

1. Introduction

1.1 Background and motivation

Accessible parking is an important component of an inclusive transportation environment. Designated spaces are commonly located near building entrances and may include additional access aisles or dimensions intended to accommodate wheelchairs, mobility devices, and assisted transfer. The availability of such spaces therefore has consequences that extend beyond the ordinary inconvenience associated with finding a parking location. When an accessible space is occupied without authorization, a person who depends on that space may be forced to park farther away, wait for another space, or abandon a trip. Effective monitoring is consequently an operational as well as an accessibility requirement.

Traditional enforcement generally depends on signage, periodic inspection, security personnel, complaints, or manual verification of permits. These approaches can be appropriate in small facilities, but they provide limited temporal coverage and can be difficult to scale across large campuses, hospitals, commercial areas, public streets, and multi-level parking facilities. A violation may occur between inspection rounds, and manual observations may not provide a consistent record of arrival time, duration, vehicle identity, or the evidence needed to support a subsequent review.

Smart-parking research has progressively introduced sensors, cameras, wireless communication, mobile applications, databases, and cloud services to improve parking availability and management. Earlier systems concentrated on occupancy, whereas newer platforms increasingly combine sensing with machine learning, computer vision, connectivity, and centralized analytics. Reviews of

smart parking have documented this transition and have highlighted the trade-offs among sensing accuracy, deployment cost, communication requirements, coverage, and information richness [3,4].

1.2 Why accessible parking is a distinct computer-vision problem

Determining whether a conventional parking space is occupied is substantially simpler than deciding whether the occupation of an accessible space is authorized. A complete enforcement decision may require answers to several related questions: Is the space designated as accessible? Is a vehicle occupying the space? Which vehicle is present? Is an accessibility permit or credential visible? Is the credential valid? Is it associated with the vehicle? Has the vehicle remained in the space for a period relevant to the local rule? Is the observed situation sufficiently clear to justify an enforcement action?

The problem can therefore be represented as a chain of evidence: accessible-space identification → vehicle detection → vehicle identity → credential detection → authorization verification → temporal verification → enforcement decision. This formulation changes the engineering objective from isolated object detection to multimodal evidence integration. A system may achieve high vehicle-detection accuracy and still produce unacceptable enforcement decisions if license-plate recognition, credential recognition, database matching, or temporal reasoning is unreliable.

1.3 Objectives and contributions

This paper provides a comprehensive technical review of AI-enabled accessible-parking monitoring and enforcement. The main objectives are to: (i) examine the evolution from conventional parking management to AI-enabled systems; (ii) analyze computer-vision techniques for parking-space, vehicle, license-plate, and accessibility-credential detection; (iii) examine the role of tracking and temporal reasoning; (iv) discuss IoT, edge, and cloud architectures; (v) compare datasets and evaluation practices; (vi) analyze privacy, security, ethical, and legal considerations; (vii) identify limitations and research gaps; and (viii) derive a conceptual evidence-centric framework for reliable and privacy-aware accessible-parking enforcement.

The principal contribution is a unified perspective that treats accessible-parking enforcement as a sequence of linked decisions rather than as a single parking-occupancy task. The review also emphasizes the distinction between technical detection confidence and enforcement confidence: a model should not automatically convert a low-confidence visual observation into a punitive action.

2. Review Scope and Approach

2.1 Scope of the review

The review covers research relevant to intelligent parking monitoring and, in particular, technologies that can contribute to accessible-parking management. The scope includes computer vision, machine learning, deep learning, object detection, semantic and instance segmentation, ALPR/ANPR, OCR, accessibility-credential recognition, multi-object tracking, IoT, edge computing, cloud platforms, databases, and evidence-based enforcement.

Because the literature uses inconsistent terminology, the conceptual scope includes accessible parking, disability parking, disabled parking, handicap parking, parking permits, accessibility badges, and related terminology. General smart-parking studies are considered when they provide enabling technologies that can reasonably support accessible-parking applications.

2.2 Literature identification

Relevant literature was identified from major scholarly sources commonly used in computer science, engineering, transportation, and smart-city research, including IEEE Xplore, ScienceDirect, SpringerLink, ACM Digital Library, Web of Science, Scopus, Google Scholar, and selected preprint repositories. Search concepts included combinations of accessible parking, disability parking, smart parking, parking occupancy, computer vision, deep learning, object detection, license-plate recognition, OCR, parking permits, IoT, edge computing, and enforcement.

Rather than treating the review as a formal systematic review, the literature is organized thematically. This approach is appropriate because the evidence spans heterogeneous study types, including application papers, datasets, computer-vision methods, system architectures, reviews, and emerging research. The objective is to synthesize the technological landscape, identify relationships among components, and critically discuss the conditions required for practical deployment.

2.3 Thematic organization

The reviewed literature is organized around the following themes: intelligent parking evolution; datasets and benchmarking; vehicle and parking-space detection; accessible-space mapping; ALPR and OCR; accessibility-credential recognition; tracking and temporal reasoning; IoT and edge/cloud architectures; evidence fusion and authorization; privacy and security; evaluation practices; and future research.

3. Evolution of Intelligent Parking Systems

Intelligent parking has evolved through several technological stages. Manual inspection represents the earliest model, followed by dedicated sensors for occupancy, camera-based monitoring, and finally integrated AI and IoT platforms. The stages should not be interpreted as strict chronological boundaries because contemporary deployments may combine several generations of technology.

Generation	Representative technologies	Primary capability	Main limitation
Parking 1.0	Manual inspection, signage	Occupancy and rule enforcement	Limited coverage; labor intensive
Parking 2.0	Ultrasonic, magnetic, infrared, RFID	Space occupancy and identification	Limited semantic information
Parking 3.0	CCTV, computer vision, ALPR	Occupancy and vehicle identity	Authorization often external
Parking 4.0	AI, IoT, edge/cloud, analytics	Detection, identification, verification, response	Privacy, uncertainty, integration

Sensor-based systems can provide reliable occupancy information without transmitting images, which may be advantageous for privacy. However, dedicated sensors generally provide limited information about vehicle identity, permit status, or the context of a violation. Camera-based systems provide substantially richer evidence because a single visual stream may contain the parking-space boundary, vehicle, license plate, credential, signage, and environmental context. The resulting information richness comes with additional challenges involving illumination, occlusion, privacy, storage, bandwidth, and model generalization.

Computer-vision parking research has often focused on classifying individual parking spaces. A dataset-centered review reported that a large proportion of published work concentrates on individual-space classification and highlighted dependence on established public datasets such as PKLot [5]. This concentration is useful for occupancy research but does not directly address the authorization problem that defines accessible-parking enforcement.

4.

Dataset	Main characteristics	Primary use	Limitation for accessible parking
PKLot	Outdoor parking imagery with labeled parking spaces	Parking-space classification	No dedicated accessibility credentials or authorization labels
CNRPark	Visual parking occupancy data	Occupancy detection	General parking context
CNRPark+EXT	About 150,000 labeled vacant/occupied patches from 164 spaces; varied cameras and conditions	Robust occupancy research	No accessible-parking enforcement labels
PLDs	Parking-lot imagery used in computer-vision studies	Benchmarking and occupancy research	Not designed for disability authorization
AccessParkCV	Aerial imagery for accessible-parking detection and space characteristics	Accessible-space mapping	Does not by itself establish vehicle authorization or violation

Datasets and Benchmarking

4.1 General parking datasets CNRPark+EXT is particularly valuable for environmental robustness because it includes variation in cameras, viewpoints, weather, illumination, shadows, and occlusions [7]. AccessParkCV addresses a complementary problem by using aerial imagery to identify accessible parking and characterize space-level properties [2]. These datasets demonstrate that the field contains useful resources for individual components, but they do not collectively provide the synchronized evidence needed for end-to-end accessible-parking enforcement.

4.2 Requirements for an accessible-parking benchmark

A purpose-built benchmark should represent the complete information chain rather than only occupancy. The dataset should contain geographically and environmentally diverse observations and should support both component-level and end-to-end evaluation.

Annotation layer	Recommended information
Scene	Camera location, timestamp, illumination, weather, viewpoint
Parking space	Polygon, accessible/non-accessible status, access aisle, occupancy
Vehicle	Bounding box, vehicle class, track identifier
Vehicle identity	Plate region, controlled or anonymized identifier
Credential	Permit/badge region, visibility, credential type, confidence
Authorization	Authorized, unauthorized, or uncertain
Temporal	Arrival, persistence, departure, duration
Enforcement	AI recommendation, human decision, reason code

Such a dataset would also enable evaluation of failure propagation. For example, a plate-recognition error may lead to an authorization error even when vehicle detection is correct. Measuring only the first-stage detector would hide this downstream effect.

5. Computer Vision for Accessible Parking

5.1 Vehicle detection

Vehicle detection has progressed from handcrafted visual descriptors and background-subtraction methods to convolutional neural networks and transformer-based detectors. One-stage detectors from the YOLO family are particularly attractive for parking applications because they can provide low-latency inference. YOLOv10, for example, was designed as an end-to-end real-time detector with an emphasis on reducing post-processing overhead [8]. Other contemporary detector families, including RT-DETR, can also be considered where accuracy, latency, and hardware constraints justify their use.

For accessible-parking enforcement, vehicle detection should not be evaluated in isolation. The detector must maintain sufficient recall under partial occlusion, crowded scenes, different vehicle sizes, night illumination, rain, glare, and camera-specific viewpoints. Small errors near parking-space boundaries can also affect the spatial association between a vehicle and an accessible space.

5.2 Parking-space occupancy

Two broad strategies dominate parking-space analysis. The first treats each predefined parking-space region as an independent classification problem. The second detects vehicles and determines occupancy by associating vehicle locations with parking-space polygons. The former can be computationally efficient when camera geometry is fixed. The latter is more flexible when vehicle-level evidence, changing occupancy patterns, or multiple vehicle classes must be considered.

For enforcement, polygon-based spatial reasoning is particularly useful because it allows the system to distinguish between a vehicle approaching a space, partially crossing the boundary, and actually occupying the designated region. The access aisle should also be represented explicitly where local rules treat it as part of the protected accessible area.

5.3 Accessible-space identification

Before a system can enforce the use of accessible parking, it must know which spaces are accessible. This information can be supplied by a maintained GIS or facility database, manually defined polygons, or computer-vision models that identify accessibility markings and signs. Aerial imagery offers an efficient way to build or update large-scale inventories. AccessParkCV demonstrates that accessible spaces can be detected and characterized from orthorectified aerial imagery [2].

A practical deployment may therefore use a two-stage architecture: an infrastructure-mapping stage creates and validates a digital inventory of accessible spaces, while a monitoring stage continuously observes those spaces. Separating these functions reduces the burden on the real-time enforcement model and allows the mapping process to be updated independently.

5.4 License-plate recognition

ALPR generally involves plate localization, geometric correction or rectification, and character recognition. Performance is influenced by illumination, camera angle, motion blur, dirt, plate design, weather, image resolution, and regional variation [9,10]. Parking applications add another challenge: the plate may be captured from a non-ideal angle because cameras are mounted for space coverage rather than dedicated plate capture.

Deep-learning detectors can localize plates in complex scenes, while OCR systems can recognize characters after cropping and rectification. Lightweight OCR architectures such as PP-OCR illustrate the continuing effort to balance recognition quality and computational efficiency [11]. For an enforcement system, exact plate-match accuracy is more meaningful than a generic OCR score because one character error can map a vehicle to the wrong authorization record.

5.5 Accessibility badge and permit recognition

Accessibility-credential recognition is a defining component of the problem. A credential may be small relative to the vehicle image, partially hidden, reflective, behind a windshield, rotated, or affected by glare. These conditions make the task closer to small-object detection under adverse imaging conditions than to ordinary object recognition.

The SHINE system is a particularly relevant example because it combines vehicle, license-plate, and disability-badge detection with centralized authorization [1]. Its reported mean average precision of 92.16% demonstrates the feasibility of integrating these visual components, but a single detection metric cannot establish complete enforcement reliability.

Future credential-recognition systems should consider both detection and authenticity/association. A visually detected badge does not necessarily establish that the credential is valid or belongs to the observed vehicle. Consequently, credential recognition should be treated as evidence that must be linked with authorization information rather than as a final decision.

5.6 Tracking and temporal reasoning

Frame-by-frame detection can generate unstable decisions. A badge may be visible in one frame and hidden in the next, while plate recognition may fluctuate because of blur or changing reflections. Multi-object tracking provides temporal continuity and allows the system to estimate arrival, persistence, departure, and parking duration.

Deep SORT combines tracking with appearance information and was designed to reduce identity switches while maintaining practical frame rates [12]. Other modern trackers, such as ByteTrack, can also be considered. For accessible parking, the important issue is not simply tracking accuracy but the ability to maintain a reliable association between a vehicle, its plate, its parking-space event, and the corresponding credential over time.

6. IoT, Edge and Cloud Architecture

IoT provides the integration layer between cameras, sensors, edge processors, communication networks, databases, dashboards, and enforcement workflows. Smart-parking architectures commonly use combinations of wired or wireless connectivity, local gateways, cloud platforms, APIs, and mobile interfaces [3,4].

Layer	Typical components	Primary function
Sensing	CCTV/IP cameras, occupancy sensors, RFID	Capture physical evidence
Edge AI	CPU/GPU/NPU gateway	Local detection, tracking, OCR

Connectivity	Ethernet, Wi-Fi, 4G/5G, MQTT	Transmit events and metadata
Platform	Database, API, rules engine	Authorization and analytics
Application	Dashboard/mobile/alert system	Human interaction
Audit	Evidence store and logs	Traceability and accountability

6.1 Edge versus cloud

Factor	Edge	Cloud	Preferred approach
Latency	Low	Usually higher	Edge for immediate alerts
Bandwidth	Low requirement	Higher for raw video	Transmit events rather than continuous video
Privacy	Potentially stronger	Greater centralized exposure	Process sensitive imagery locally where feasible
Compute	Limited	High	Hybrid
Central analytics	Limited	Strong	Cloud for aggregated metadata
Offline operation	Possible	Limited	Edge fallback

A hybrid architecture is therefore attractive. Raw imagery can be processed locally, while only the minimum evidence required for authorization and human review is transmitted to a central service. This can reduce bandwidth and exposure of personally identifiable information while retaining centralized analytics for system-level management.

7. AI-Based Accessible-Parking Enforcement

7.1 From occupancy detection to authorization

The central distinction between ordinary smart parking and accessible-parking enforcement is authorization. Occupancy answers the question 'Is something in the space?' Enforcement must answer 'Is this particular vehicle authorized to use this particular accessible space under the applicable rules at this time?'

Accordingly, an intelligent enforcement pipeline should integrate spatial evidence, visual identity, credential information, database information, and temporal context. A possible functional chain is:

Accessible-space map → vehicle detection → tracking → plate detection/OCR → credential detection/OCR → authorization database → temporal/spatial reasoning → confidence estimation → human review or warning → auditable record.

7.2 Authorization as a policy-engine problem

Authorization should not be reduced to a binary visual classification. Permit validity, vehicle association, time restrictions, space-specific rules, temporary permissions, and jurisdictional requirements may differ across locations. Therefore, the AI layer should provide evidence to a policy engine rather than independently determine legal entitlement.

A robust authorization service may return three states: Authorized, Unauthorized, and Uncertain. The uncertain state is essential when visual evidence is incomplete, database information is unavailable, or confidence falls below a predefined threshold. Such cases should be routed to a trained human reviewer instead of being converted automatically into a punitive action.

7.3 Evidence fusion

Accessible-parking enforcement requires evidence from multiple sources. Let V represent vehicle-detection confidence, P plate-recognition confidence, B credential-detection confidence, A authorization confidence, S spatial-association confidence, and T temporal-persistence confidence. A simple conceptual fusion score can be expressed as:

$$C = w_v V + w_p P + w_b B + w_a A + w_s S + w_t T, \quad \text{where } \sum w = 1$$

This equation is a conceptual formulation rather than a validated enforcement model. In a deployed system, the weights and thresholds should be learned or calibrated using representative validation data. Alternative approaches such as probabilistic graphical models, Bayesian fusion, Dempster–Shafer reasoning, or learned multimodal fusion may be evaluated.

The fusion output should support calibrated decisions such as authorized, unauthorized with sufficient evidence, or uncertain. The system should preserve the individual evidence components so that a reviewer can understand why an alert was generated.

8. Comparative Analysis of Technologies

Technology	Strength	Key limitation	Role in accessible parking
CCTV	Rich visual context	Lighting, bandwidth, privacy	High
Ultrasonic/magnetic sensors	Simple occupancy measurement	Limited semantic information	Medium
Vehicle detection	Real-time localization	Dataset and environmental dependence	High
ALPR	Vehicle identity	Blur, angle, privacy	Very high
OCR	Text recognition	Small/blurred text	High
Credential detection	Direct accessibility evidence	Small, occluded, reflective objects	Very high
Tracking	Temporal consistency	Identity switches	Very high
RFID/database linkage	Explicit identity/authorization	Infrastructure dependence	High
Edge AI	Low latency and local processing	Hardware constraints	Very high
Cloud AI	Centralized analytics and scalable compute	Latency and privacy exposure	High
Aerial imagery	Large-scale space mapping	Limited temporal enforcement evidence	Medium

9. Evaluation Metrics and Experimental Protocols

Evaluation should be performed at both component and system levels. Reporting only detector accuracy is insufficient because an enforcement platform can fail even when individual components perform well.

Subsystem	Recommended metrics
Vehicle/space detection	Precision, recall, mAP, IoU
Occupancy	Accuracy, balanced accuracy, F1, false-occupancy rate
ALPR/OCR	Character accuracy, exact plate-match rate, CER/WER
Credential detection	Precision, recall, mAP, small-object recall
Tracking	IDF1, MOTA, identity switches, track fragmentation
Violation detection	Precision, recall, F1, false-positive rate
End-to-end system	Latency, event latency, FPS, uptime
Edge deployment	Power, memory, model size, throughput
Human review	Review time, agreement, overturn rate

Cross-condition evaluation is particularly important. Models should be tested across camera locations, day and night, different weather conditions, viewing angles, vehicle types, occlusion levels, and geographic plate formats. Cross-site testing is essential because performance on a familiar dataset may not represent performance after a change in camera geometry, parking layout, illumination, or regional credential design [5,7].

An additional requirement is end-to-end error analysis. For example, researchers should report how often a correct vehicle detection is converted into an incorrect authorization decision because of plate recognition, credential recognition, database matching, or temporal-association errors.

10. Privacy, Security, Ethical and Legal Considerations

ALPR can transform a vehicle plate into a persistent identifier associated with time and location. Consequently, an accessible-parking monitoring system may create a sensitive mobility record even when its original purpose is limited to parking enforcement. Privacy should therefore be incorporated into the architecture from the beginning rather than added after deployment.

Risk	Example	Potential mitigation
False positive	Authorized vehicle flagged	Confidence thresholds and human review
False negative	Unauthorized occupation missed	Temporal fusion and multiple views
Privacy	Persistent plate-location history	Data minimization and retention limits
Security	Unauthorized access to records	Encryption, authentication, role-based access, audit logs
Model bias	Performance varies across cameras, vehicles or plate formats	Diverse validation and subgroup analysis
Function creep	Parking infrastructure becomes general surveillance	Purpose limitation and governance
Credential misuse	Plate and permit do not correspond	Database linkage and uncertainty state

Privacy-preserving design can include local inference, selective transmission, plate tokenization or controlled identifiers, encryption in transit and at rest, role-based access control, short retention periods, and automated deletion of non-event imagery. Federated learning may also support multi-site model improvement without requiring all raw images to be centralized, although its suitability depends on the deployment context.

Legal requirements for accessible permits and enforcement vary across jurisdictions. An AI system should therefore be configurable through a policy layer rather than embedding legal assumptions directly into the vision model. The technical system should assist authorized personnel and should not be treated as an autonomous legal authority.

11. Proposed Evidence-Centric Framework

The reviewed literature supports a modular architecture in which each stage contributes evidence to the next stage. The proposed framework is conceptual and is intended to organize future implementation and evaluation rather than to claim experimental validation.

Module	Candidate technologies	Priority
Accessible-space mapping	GIS, segmentation/detection, aerial imagery	Very high
Vehicle detection	YOLO-family, RT-DETR or comparable real-time detector	High
Tracking	Deep SORT, ByteTrack or comparable tracker	High
Plate detection	YOLO-family or transformer detector	High
OCR	PP-OCR or plate-specific recognizer	High
Credential detection	Custom detector, YOLO/RT-DETR/transformer	Very high
Authorization	Secure API/database and policy engine	Very high
Evidence fusion	Rules, probabilistic fusion or learned multimodal model	Very high
Edge deployment	GPU/NPU/CPU gateway	High
Human review	Web/mobile dashboard with evidence packet	Very high

The system should produce an evidence packet for each candidate event. The packet may contain an event identifier, accessible-space identifier, time stamp, vehicle image, plate evidence, credential evidence, authorization response, parking duration, confidence values, and the reason for the recommendation. Such an architecture improves auditability and allows human reviewers to challenge or overturn an AI recommendation.

12. Research Gaps and Open Challenges

Research gap	Current situation	Research opportunity

Dedicated datasets	General parking datasets dominate	Create multimodal accessible-parking benchmarks
End-to-end evaluation	Most studies isolate individual tasks	Evaluate space → vehicle → plate → credential → authorization
Credential recognition	Limited direct research	Develop small-object and occlusion-robust models
Environmental robustness	Lighting, weather and camera variation affect performance	Use day/night/weather/cross-camera benchmarks
Cross-site generalization	Dataset dependence is a recurring concern	Perform geographic and cross-location testing
Temporal reasoning	Frame-level metrics may miss parking behavior	Use track-based violation protocols
Privacy	ALPR can create persistent mobility records	Use edge processing, minimization and retention controls
Uncertainty	Binary decisions can be unsafe	Use calibrated three-state decisions
Standardization	Metrics and datasets differ	Develop common reporting and benchmark protocols
Human factors	Technical performance may overlook user needs	Evaluate human review, accessibility outcomes and trust

13. Failure Modes Requiring Explicit Evaluation

An important direction for future research is systematic failure analysis. Several situations can cause a technically plausible system to generate an incorrect enforcement recommendation.

Failure scenario	Potential cause	Required response
Credential not visible	Occlusion or windshield position	Do not infer unauthorized status solely from absence
Credential visible but unreadable	Blur, glare or low resolution	Use temporal evidence or human review
Credential belongs to another vehicle	Misuse or database mismatch	Cross-check authorization and vehicle identity
Plate recognition error	Angle, blur, dirt or glare	Use multiple frames/views and confidence calibration
Temporary stopping	Vehicle enters the polygon briefly	Use temporal persistence and local rules
Partial occupancy	Vehicle crosses a boundary but does not fully park	Use spatial and temporal reasoning
Camera obstruction	People, vehicles or infrastructure block view	Mark event uncertain
Night-time failure	Low illumination and glare	Use appropriate illumination/sensors and night-specific validation

Incorrect space map	Outdated or inaccurate GIS information	Maintain and periodically validate the accessible-space inventory
---------------------	--	---

14. Future Research Directions

Future research should move from isolated demonstrations toward robust, reproducible and human-centered systems. The following directions are particularly important.

- Develop an open or responsibly governed accessible-parking dataset containing synchronized space, vehicle, plate, credential, tracking, authorization, and enforcement annotations.
- Evaluate models across cameras, sites, weather, time of day, parking geometries, vehicle types, and geographic regions.
- Develop lightweight credential detectors optimized for small objects behind windshields and under reflections or partial occlusion.
- Combine RGB vision with depth, thermal, radar, RFID, or occupancy sensors when the additional information provides measurable benefit.
- Use edge AI to reduce raw-video transmission, latency, and privacy exposure.
- Develop calibrated uncertainty models and an Authorized/Unauthorized/Uncertain decision policy.
- Generate explainable evidence packets that preserve the visual and database evidence supporting each alert.
- Investigate federated and other privacy-preserving learning approaches for multi-site model improvement.
- Develop standardized end-to-end metrics for accessible-parking enforcement, including false-positive enforcement rate and human overturn rate.
- Evaluate human review time, agreement, accessibility outcomes, user trust, and operational cost in addition to conventional model metrics.

15. Limitations of the Review

This paper is a comprehensive technical review rather than a formal systematic review or quantitative meta-analysis. The literature is interdisciplinary and terminology is inconsistent; relevant work may use terms such as accessible parking, disability parking, parking permits, badges, smart parking, or parking enforcement. The reviewed studies also use heterogeneous datasets, metrics, hardware, environmental conditions, and experimental protocols, making direct numerical comparison difficult.

The field is evolving rapidly, and recent technical developments may initially appear in conference proceedings or preprint repositories before journal publication. In addition, legal requirements for accessible permits and enforcement differ across jurisdictions. Accordingly, the proposed architecture should be interpreted as a research framework and not as a universally valid legal or operational specification.

16. Conclusion

AI, computer vision, ALPR, OCR, IoT, and edge computing provide the technical foundations for continuous accessible-parking monitoring. The literature demonstrates substantial maturity in general parking occupancy detection, vehicle detection, license-plate recognition, tracking, and smart-parking architectures. However, accessible-parking enforcement introduces a more demanding requirement: the system must establish whether a particular vehicle is authorized to occupy a particular accessible space at a particular time.

The review therefore identifies accessible-parking enforcement as an evidence-integration problem rather than a single object-detection task. Accessible-space mapping, vehicle detection, tracking, plate recognition, credential analysis, authorization databases, temporal reasoning, and human review must operate as coordinated components. Existing datasets are not yet sufficient for this end-to-end task, and future benchmarks should include synchronized annotations covering the full decision chain.

A practical architecture should also treat privacy, security, uncertainty, explainability, and human oversight as core requirements. The most appropriate operational model is not necessarily AI-to-fine automation; instead, AI should generate structured evidence and confidence estimates, while uncertain cases are reviewed by authorized personnel. Future research should consequently focus on cross-condition validation, multimodal datasets, privacy-preserving edge intelligence, calibrated evidence fusion, and accessibility-centered evaluation. Such developments can move smart parking beyond occupancy management toward responsible, reliable, and inclusive protection of accessible parking.

References

- [1] Neupane, D., Bhattarai, A., Aryal, S., Bouadjene, M. R., Seok, U.-M., & Seok, J. (2024). SHINE: A deep learning-based accessible parking management system. *Expert Systems with Applications*, 238, pp. 122205. <https://doi.org/10.1016/j.eswa.2023.122205>.
- [2] Hwang, J., Li, C., Kang, H., Hosseini, M. and Froehlich, J.E., (2025). "Where Can I Park?" Understanding Human Perspectives and Scalably Detecting Disability Parking from Aerial Imagery. In *Proceedings of the 27th International ACM SIGACCESS Conference on Computers and Accessibility* (pp. 1-20).
- [3] Biyik, C., Allam, Z., Pieri, G., Moroni, D., O'Fraifer, M., O'Connell, E., Olariu, S., & Khalid, M. (2021). Smart Parking Systems: Reviewing the Literature, Architecture and Ways Forward. *Smart Cities*, 4(2), 623–642. <https://doi.org/10.3390/smartcities4020032>.
- [4] Khanna, A., & Anand, R. (2021). Smart parking systems: comprehensive review based on various aspects. *Heliyon*, 7(5), e07050. <https://doi.org/10.1016/j.heliyon.2021.e07050>.
- [5] Almeida, P. R. L. de, Alves, J. H., Parpinelli, R. S., & Barddal, J. P. (2022). A systematic review on computer vision-based parking lot management applied on public datasets. *Expert Systems with Applications*, 198, 116731. <https://doi.org/10.1016/j.eswa.2022.116731>.
- [6] Amato, G., Carrara, F., Falchi, F., Gennaro, C., & Vairo, C. (2015). CNRPark+EXT: A dataset for visual occupancy detection of parking lots. Dataset resource, CNR-EXT/CNRPark.
- [7] Wang, A., Chen, H., Liu, L., Chen, K., Lin, Z., Han, J. and Ding, G., (2024). Yolov10: Real-time end-to-end object detection. *Advances in neural information processing systems*, 37, pp.107984-108011.
- [8] Shashirangana, J., Padmasiri, H., Meedeniya, D. and Perera, C., (2020). Automated license plate recognition: a survey on methods and techniques. *IEEE Access*, 9, pp.11203-11225. <https://doi.org/10.1109/ACCESS.2020.3047929>.
- [9] Saha, S. (2019). A Review on Automatic License Plate Recognition System. *arXiv:1902.09385*.
- [10] Du, Y., Li, C., Guo, R., Yin, X., Liu, W., Zhou, J., Bai, Y., Yu, Z., Yang, Y., Dang, Q., & Wang, Y. (2020). PP-OCR: A Practical Ultra Lightweight OCR System. *arXiv:2009.09941*.
- [11] Wojke, N., Bewley, A., & Paulus, D. (2017). Simple Online and Realtime Tracking with a Deep Association Metric. *IEEE International Conference on Image Processing (ICIP)*, 3645–3649. <https://doi.org/10.1109/ICIP.2017.8296962>.
- [12] De Almeida, P. R. L., Oliveira, L. S., Britto, A. S., Silva, E. J., & Koerich, A. L. (2015). PKLot – A robust dataset for parking lot classification. *Expert Systems with Applications*, 42(11), 4937–4949. <https://doi.org/10.1016/j.eswa.2015.02.009>.
- [13] Laroca, R., Zanlorensi, L.A., Gonçalves, G.R., Todt, E., Schwartz, W.R. and Menotti, D., (2021). An efficient and layout-independent automatic license plate recognition system based on the YOLO detector. *IET Intelligent Transport Systems*, 15(4), pp.483-503.
- [14] Wang, Y., Bian, Z.P., Zhou, Y. and Chau, L.P., (2021). Rethinking and designing a high-performing automatic license plate recognition approach. *IEEE transactions on intelligent transportation systems*, 23(7), pp.8868-8880.
- [15] Li, C., et al. (2022). PP-OCRv3: More Attempts for the Improvement of Ultra Lightweight OCR System. *arXiv:2206.03001*.
- [16] Terven, J., Córdova-Esparza, D.M. and Romero-González, J.A., (2023). A comprehensive review of yolo architectures in computer vision: From yolov1 to yolov8 and yolo-nas. *Machine learning and knowledge extraction*, 5(4), pp.1680-1716.