

Global Journal of Engineering and Technology Review

Volume: 02 Issue: 09 September 2026

ParkIQ A Predictive Framework for Urban Parking and Accountable Kerbside Management

Syed Darda Rahman

School of Computer Science and Technology, University of Bedfordshire Luton, United Kingdom

Published Date: September 30, 2026

Article DOI : 10.65150/EP-gjetr/V2E9/2026-16

ABSTRACT: Parking information is often split between occupancy sensors, payment services and local restriction records. This makes it difficult to tell drivers where they may park, estimate whether a space will remain available, and explain an enforcement decision from a consistent record. This paper proposes ParkIQ, an engineering framework that connects these tasks through a shared representation of individual parking bays. The design combines six layers: sensing, a geospatial bay model, time-dependent restrictions, data fusion, occupancy forecasting and user services. A temporal convolutional network and a gradient-boosted model are proposed for probabilistic forecasts at 15-minute intervals up to 90 minutes ahead. An associated enforcement workflow preserves evidence for authorised review, while a separate commercial model supports bookings and service provision. The contribution is a system specification and prospective evaluation protocol; no trained-model results, operational savings or deployment outcomes are reported. The protocol compares the proposed ensemble with simpler baselines, separates training from future observations, and tests calibration, sensor failures and transfer between locations. It also distinguishes evidence integrity from the correctness and legal validity of an enforcement decision. The framework provides a testable basis for investigating whether shared bay information can improve parking guidance and operational accountability. Its value will depend on data quality, lawful access to information, user participation and evidence from a controlled pilot.

KEYWORDS: smart parking, occupancy forecasting, kerbside management, sensor fusion, accountable enforcement, urban infrastructure

1. INTRODUCTION

A driver approaching a destination needs to know more than whether a parking space is empty. The space must be suitable for the vehicle, available at the expected arrival time and permitted for the intended duration. Payment applications, occupancy feeds and restriction records can each answer part of this question. The engineering difficulty is to keep their answers consistent as vehicles arrive, restrictions change and observations become stale.

Parking search can add unnecessary circulation to busy streets. Shoup's (2006) analysis describes this mechanism across different urban settings, but its findings should not be read as a fixed percentage of traffic in every city. A UK study published by INRIX in 2017 estimated 44 hours of parking search per driver annually and a national cost of £23.3 billion (INRIX, 2017). These historical estimates establish the scale of the problem at the time of the study; they are not measurements of present-day conditions or estimates of what ParkIQ would save.

Parking also involves competing uses of limited street space. Loading, accessible parking, residents' permits, short stays and electric-vehicle charging can impose different requirements on the same location. Sha et al. (2024) examined alternative uses of on-street parking in a model of Leicester, illustrating how changes in kerbside allocation can affect traffic, safety and emissions under specific connected-vehicle scenarios. Such scenario results motivate careful evaluation; they do not establish the effect of a new parking platform.

ParkIQ is proposed as a shared information layer for drivers, authorities and space operators. Its principal design choice is to connect sensing, restrictions, prediction and evidence through a versioned bay record. This paper specifies the architecture, the forecasting problem, an evidence-preserving enforcement workflow and a prospective evaluation. Commercial arrangements are examined as deployment assumptions rather than demonstrated sources of income. The contribution lies in making these dependencies explicit and testable. It does not claim that integration is absent from all existing products or that the proposed ensemble outperforms established models.

2. LITERATURE REVIEW

2.1 Sensing and service integration

Smart parking has been studied as a combination of information collection, infrastructure deployment and services for users. Lin et al. (2017) organise this broad ecosystem into those interconnected themes. Barriga et al. (2019) examine sensors, communication protocols and software components, showing why selecting a detector alone is insufficient to specify a working service. Fahim et al. (2021) compare systems across their sensing, networking, interfaces and computational approaches. Biyik et al. (2021) likewise discuss the relationship between detection technologies and communication infrastructure.

More recent reviews cover context-aware urban systems and the digitalisation of parking services (Zulfiqar et al., 2023; Channamallu et al., 2026). These studies support an integration-oriented design, while also showing that integration itself is not a sufficient novelty claim. Rocco et al. (2023) consider an architecture serving public authorities and private users, including reservation-based parking. ParkIQ therefore needs to be assessed against comparable service architectures, not only against standalone prediction algorithms.

The observation layer introduces practical constraints. Perković et al. (2020) examine sensor performance and energy use, which are relevant to maintenance and battery replacement in a dispersed installation. Zeng et al. (2024) place parking sensors within a broader review of IoT applications, communication and data processing in smart cities. Navpreet et al. (2025) review parking-space detection and classification methods and their evaluation datasets. Together, these studies motivate explicit treatment of missing readings, occlusion, device health and the distinction between a sensor observation and independently checked ground truth.

2.2 Parking prediction

Vlahogianni et al. (2016) investigate real-time parking prediction using observations from an instrumented city. Xiao et al. (2018) develop a model-based approach to parking availability, while Yang et al. (2019) combine multiple sources of spatial and temporal information in a deep-learning framework. These are useful precedents, but they address different targets and operating conditions. A model predicting a car park's available capacity does not necessarily solve the question of whether one particular bay will be free when a driver arrives.

Recent surveys make these distinctions central to the choice of data, models and metrics (Ma et al., 2024; Xiao et al., 2023; Cao et al., 2024). Representation learning for parking data management is considered by Yang et al. (2022). Liu et al. (2021) study regional inflow and outflow for shared parking, and Feng et al. (2022) model spatial and temporal relationships in zone-level availability. Hao et al. (2023) connect spatiotemporal learning to curb management. These studies suggest that both temporal history and relationships between locations deserve attention, although neither should be assumed to improve every deployment.

Operational constraints also matter when choosing a model. Dia et al. (2022) compare predictive performance with execution time and resource requirements. Enríquez et al. (2024) combine occupancy prediction with a fog-computing architecture and examine response delay. ParkIQ accordingly treats a temporal convolutional network and gradient boosting as candidates to test. A more complex ensemble would be justified only if it improves useful forecast quality without unacceptable computational or maintenance costs.

2.3 Privacy and accountable decisions

Al Amiri et al. (2019) propose a parking system using a consortium blockchain and private information retrieval to protect the privacy of parking queries and reservations. That use case differs from preserving evidence for an enforcement decision. ParkIQ considers a ledger only as an optional integrity mechanism for evidence records. A hash can help reveal later alteration; it cannot establish that a camera identified the right vehicle, a sensor observed the right bay, or a restriction was correctly applied.

The proposed evaluation also draws on work outside parking. Roberts et al. (2017) explain why validation must respect temporal and spatial dependence. Kapoor and Narayanan (2023) show how leakage can produce optimistic machine-learning findings. Gneiting et al. (2007) provide a foundation for evaluating probabilistic forecasts through calibration and sharpness, and Gama et al. (2014) review adaptation to concept drift. These principles inform the test protocol rather than supplying evidence of ParkIQ's performance.

3. MATERIALS AND METHODS

3.1 Study design and scope

This is a design-and-evaluation protocol. The research artefact is a proposed architecture with data contracts, decision rules and a prospective validation plan. No pilot dataset was collected or model trained for this paper. The literature review is a focused narrative synthesis of relevant engineering work, not a systematic review or meta-analysis. Its purpose is to identify design choices and tests, rather than to estimate a pooled effect or establish an exhaustive novelty claim.

The intended setting is urban parking in the United Kingdom. The initial enforcement discussion is scoped to England because the cited statutory guidance applies there (Department for Transport, 2022). Deployment elsewhere would require a separate review of the applicable rules. Likewise, public civil enforcement and private parking management are separate operating regimes; a shared software platform must not treat their powers, notices or appeal procedures as interchangeable.

Three user groups define the requirements. Drivers need understandable restrictions and credible availability information. Authorities need reliable observations, auditable decisions and tools for reviewing exceptions. Space operators need controlled listing, booking and settlement. Table 1 connects these needs to proposed functions and observable tests. All functions remain design proposals.

Table 1. User requirements and proposed tests

User need	Proposed function	Evidence needed
Find a suitable space	Arrival-time availability with restrictions	Forecast quality; suitability errors; observed search time
Explain a possible contravention	Versioned rules and retrievable evidence	Independent case audit; incorrect-case rate; review time
Manage private capacity	Listing, booking and settlement	Realised occupancy; cancellations; net operating margin
Manage changing bay uses	Time-dependent rules and charger status	Rule-update delay; stale-status rate; exception handling
Maintain service reliability	Freshness checks, fallback and rollback	Outage tests; latency; recovery time

3.2 Bay records and data provenance

Each bay would have a stable identifier, geometry, ownership, bay type, sensor bindings and a history of applicable restrictions. A restriction record would include its source, version, effective dates, permitted uses, maximum stay and relevant exceptions. A prediction or enforcement case would refer to the rule version that applied when the event occurred, allowing later reviewers to reconstruct the decision.

Ordnance Survey geometry and authority records are candidate inputs, subject to licensing, coverage and local data quality. Neither mapping geometry nor a payment record is assumed to provide a complete, authoritative bay inventory. Before a pilot, the proposed inventory would be reconciled with field observations, signage and the authority's restriction records. Unresolved locations would remain unavailable for automated case progression.

Events would carry both the time of observation and the time they reached the platform. This distinction matters when a delayed reading arrives after a forecast has been issued. Sensor health, source identity and confidence would be retained alongside the observation. Conflicting or stale observations would produce an uncertain state rather than a forced occupied-or-free label. Sensitive identity data would be separated from routine occupancy features.

3.3 Six-layer architecture

Figure 1 presents the six-layer design. The sensing layer receives occupancy transitions from candidate magnetometer or radar devices, camera observations where appropriate, and optional aggregated app signals. LoRaWAN and NB-IoT are candidate communications technologies. A three-second sensor reporting interval is a design target to test, not a measured device capability. Camera observations at a zone entrance cannot by themselves identify the bay occupied by a vehicle; any such association would need a separately validated method.

The second layer maintains the geospatial bay record. The third applies time-dependent restrictions and exceptions. The fourth reconciles observations with their provenance and freshness. These layers produce an operational state that keeps physical occupancy, permitted use and payment status distinct. For example, a vehicle can occupy a bay while the platform lacks sufficient evidence to decide whether its use is permitted. The fifth layer issues occupancy forecasts. The sixth exposes appropriately limited information to a driver application, authority dashboard and operator portal through a common API. Proposed driver functions include availability guidance, restriction explanations, payment and reminders at 80% and 95% of the permitted stay. These reminders are service features, not replacements for signage or legal notice requirements. Authority functions include review queues and dispatch support; operator functions include listing, availability management and settlement.




Figure 1. Proposed six-layer architecture. All layers refer to a versioned bay record; the diagram represents a design rather than a deployed system.

3.4 Candidate data inputs

Table 2 specifies 18 input categories and separates forecast inputs, operational records and evaluation outputs. This is a candidate data specification, not an acquired training corpus. Each source would need an owner, access agreement, coverage assessment and quality check before use.

In particular, prediction errors become available only after the relevant outcome has been observed. They can support monitoring and later retraining but cannot be features for the prediction being scored. Emission factors belong in the impact assessment, not in the occupancy target. App interactions measure behaviour among participating users and may not represent all drivers. Opt-in GPS signals would require a documented collection purpose and safeguards before entering any analytical pipeline.

Table 2. Candidate data specification

Category	Candidate source / cadence	Intended use and constraint
1 Sensor occupancy	In-ground IoT / 3 s target	Observation and candidate label; independent audit required
2 ANPR observations	Authorised cameras / event-driven	Identity and entry/exit evidence; does not itself locate a bay
3 GPS density	Opt-in app / real time	Aggregated demand context; participation bias and privacy controls
4 Bay geometry	Licensed mapping / nightly reconciliation	Location and zone features; field validation needed
5 Parking restrictions	Authority records / change-driven	Rule interpretation; preserve versions and urgent updates
6 Session history	Platform records / continuous	Past use patterns; future duration excluded from features
7 Prediction errors	Evaluation log / after outcome	Monitoring and later retraining only
8 Weather	Forecast provider / hourly	Forecast available at issue time; no realised future weather
9 Event schedules	Local or organiser feeds / weekly and updates	Known demand context; preserve publication timestamps
10 School term dates	Local authority or school calendar / termly	Local calendar effects; check school-specific exceptions
11 Road closures	Relevant road authority / updates	Known access changes; distinguish scheduled and observed events
12 Traffic flow	Local or national feeds / 15 min target	Arrival-time context; record feed delays
13 Charger telemetry	Charge-point operator / real time	Charger status, separately from parking occupancy
14 Transport timetables	Rail and bus feeds / daily and updates	Scheduled connections and park-and-ride context
15 Cross-city patterns	Permitted aggregates / nightly candidate	Transfer experiments; disclosure-risk and domain-shift checks
16 Emission factors	Applicable official factors / annual	Impact assessment only; not occupancy labels
17 Utilisation history	Occupancy records / continuous	Past utilisation; derive within training periods
18 Interaction logs	Permitted app telemetry / continuous	User-service evaluation; not representative of all drivers

3.5 Prediction target and proposed ensemble

For bay b , forecast issue time t and horizon h , let $y(b,t+h)$ equal 1 when the bay is physically occupied and 0 when it is free. Missing or disputed observations are excluded from the primary label set and reported separately. Horizons are 15, 30, 45, 60, 75 and 90 minutes. The model estimates the probability of occupancy using information available at the issue time:

$$p(b,t,h) = P[y(b,t+h) = 1 \mid \text{information available at time } t] \quad (1)$$

The first candidate model is a temporal convolutional network using recent occupancy sequences and time indicators. The second is a gradient-boosted model, with XGBoost proposed as an implementation option, using lagged occupancy and contextual features such as weather forecasts, scheduled events, road closures and observation quality. Features must reflect what was known at time t . Realised future weather, completed session duration and later corrections to a rule record must not enter a retrospective forecast as if they were available earlier.

For each horizon, a weighted combination of the two predictions would be fitted on development data. The combination weight and any calibration transformation would be selected without using the final test period. The architecture does not fix an optimal network depth, history length or tree configuration before data are available. A pilot protocol would record the search space, optimisation budget, random seeds and selected settings before final testing.

An occupancy probability is not a direct measure of sensor reliability or legal confidence. The interface would display predicted availability, observation freshness and data-quality status separately. For initial testing, predictions with $\max(p, 1 - p)$ below 0.65 would be labelled uncertain. This decision rule would be validated separately from calibration and data quality. Uncertain forecasts should remain visible with a clear qualification, and guidance should fall back to coarser zone-level information where individual bays cannot be observed reliably.

3.6 Validation and model updating

The evaluation would use chronological training, validation and test periods, with a gap of at least the maximum forecast horizon where necessary to prevent target overlap. All preprocessing, feature selection, ensemble weights and calibration would be fitted within the development portion. Nearby observations and repeated records from the same parking event would not be randomly distributed across train and test sets. Separate tests

would hold out locations to assess spatial transfer, following the distinction between interpolation and extrapolation discussed by Roberts et al. (2017).

The proposed update cycle is shown in Figure 2. A trailing 90-day window and threefold weighting of the latest seven days are initial choices for comparison with unweighted and alternative-window baselines. The final 24 hours would be reserved for candidate validation, outside the training window; labels must have matured before evaluation. A single day may be unrepresentative, so promotion would also require acceptable performance across longer rolling validation periods and relevant operating conditions.

Retraining is provisionally scheduled for 02:00 local time, with an explicit policy for daylight-saving transitions and UTC event storage. The existing model would remain active unless the candidate passes predeclared forecast, calibration and service checks. Persistent sensor errors would trigger investigation rather than automatic repeated fitting to corrupted labels. Versioned models, rollback and drift monitoring would make the update process inspectable (Gama et al., 2014). No claim is made that this process has been implemented or timed.




Figure 2. Proposed model update cycle. Completed outcomes support later updates, while final test data remain excluded from model selection.


3.7 Evidence-preserving enforcement workflow

Figure 3 shows a proposed sequence from an observation to a reviewed case. A possible contravention would first be checked against the applicable rule, vehicle-to-bay association and available exemptions. Missing evidence or inconsistent records would route the case to review or closure. A forecast of future occupancy would never be used as evidence that a contravention occurred.

The evidence package would record timestamps, source identifiers, the relevant rule version, observation quality and authorised review actions. Where justified, encrypted images or other evidence would be held in access-controlled storage. An optional permissioned ledger would store integrity references rather than raw vehicle identifiers or images. Even a digest can remain linkable in some contexts, so its retention and access controls would need assessment alongside the underlying record.

An observation-confidence threshold of 0.85 would be tested as a candidate rule for case triage. It would not, by itself, authorise a penalty. Notification, applicable grace periods, keeper-data access and notice issuance would follow the relevant authority's lawful procedure. England's civil parking enforcement guidance limits the circumstances in which camera evidence can support postal notices; it does not provide a general route to automatic penalties for every sensed bay (Department for Transport, 2022). A configurable timer cannot override those restrictions.

Human review would be the default pilot arrangement. The proposed 48-hour internal override window is an operational service target, not a limit on statutory challenge or appeal rights. Driver reminders would be optional assistance: failure to receive an app alert would not be treated as evidence against a motorist. Retention would follow documented purpose and necessity, rather than a blanket seven-year rule. ICO guidance on ANPR supports minimising retention of information about vehicles that are no longer relevant to the processing purpose (Information Commissioner's Office, n.d.).



At every stage: uncertain or insufficient evidence can lead to review or closure.

Figure 3. Proposed evidence and review workflow. Uncertainty can stop case progression; an integrity log does not establish the correctness or legality of a decision.

3.8 Commercial assumptions and cost assessment

Table 3 records the proposed indicative commercial ranges. They are assumptions for sensitivity analysis, not market quotations, contracts or forecasts. The assessment would distinguish subscription income and booking revenue from penalty receipts. Civil enforcement should pursue compliance and transport objectives rather than revenue or notice-volume targets (Department for Transport, 2022). A per-case processing charge would therefore need scrutiny for incentives as well as cost recovery.

RAC Foundation's analysis reported £2.163 billion of parking income and a £1.094 billion surplus among 295 English councils reporting for 2024-25 (RAC Foundation, 2025). These totals indicate the financial scale of parking operations; they are not an addressable-revenue estimate for ParkIQ. No share of that income can be attributed to the proposed system without local evidence.

A simple private-site example illustrates the importance of explicit assumptions. With 200 spaces, 30% utilisation, one £3 paid session per occupied space per night and 30 nights, gross monthly receipts would be £5,400. At a 5-8% commission, the platform's gross share would be £270-£432 before costs. The calculation assumes demand and permission to offer the spaces; neither has been demonstrated. It also excludes payment charges, tax, refunds, staffing, installation and maintenance.

Any pilot business case would include sensor replacement, connectivity, mapping licences, cloud services, support, officer review, appeals and cybersecurity. The framework does not assume that detection technology replaces a specified number of officers. Carbon-credit issuance remains a speculative future option requiring a recognised methodology and independent verification; an estimated reduction in driving is not automatically a tradable credit.

Table 3. Illustrative commercial assumptions

Stream	Charging basis	Proposed range or condition
Council service licence	Per city per month	£8,000-£15,000
Case processing	Per notice / case	£1.50-£3.00; incentive and procurement review required
Business subscription	Per month	£99-£999
Booking commission	Share of booking value	5-8%; depends on actual booking price
Data API licence	Per city feed per year	£10,000-£50,000; subject to source rights
Carbon credits	Potential future activity	No forecast value; methodology and verification unresolved
White-label licence	Per year	£100,000-£500,000; no contracts assumed

3.9 Prospective evaluation protocol

The initial deployment plan specifies a single-city pilot of 200-500 sensors and 3-5 candidate ANPR locations, subject to suitability and approval. Months 1-6 would prioritise bay mapping, data quality, guidance and shadow-mode case review. Months 7-12 would test forecasting and, only where authorised, additional operational functions. Months 13-18 would examine pricing, impact accounting and external data interfaces. Expansion to five further cities during months 19-30 would depend on the preceding evidence. These are planning stages, not confirmed partnerships or a funded timetable.

Before collection begins, the study team would register the final hypotheses, outcome definitions, exclusion rules and analysis plan. This paper does not claim that such registration has already occurred. Sample size and duration would be justified from expected event rates and clustering by location and time, rather than from sensor count alone. Independent manual audits would estimate label error and examine disagreements between observation sources.

Table 4 specifies five evaluation questions. Forecast comparisons would include persistence, time-of-week frequency, logistic regression, gradient boosting alone and the temporal network alone. A spatial model would be included if the available network supports a meaningful comparison. All models would use the same information cut-off, eligible observations and testing periods. Ablations would remove contextual features, recency weighting, cross-location transfer and individual sensor sources in turn.

Brier score and log loss would assess probabilistic forecasts, with reliability plots and calibration error by horizon. Precision, recall and class prevalence would accompany any threshold-based accuracy result. Accuracy aspirations of 70% by month 6 and 82% by month 12 are exploratory planning values rather than primary success criteria: a highly occupied area could exceed them with an unhelpful constant prediction. Uncertainty intervals would account for temporal and spatial clustering.

For operational outcomes, the preferred design is a phased rollout with comparable untreated locations and a pre-intervention baseline. Analysis would account for changes in tariffs, restrictions, traffic demand and observation coverage. Detecting more contraventions after installing sensors is not necessarily an increase in actual violations. A common audit method is therefore needed before and after implementation. Claims about search time or emissions would require observed travel outcomes and a stated counterfactual, not only app activity or model accuracy.

Table 4. Prospective evaluation questions

Question	Primary comparison or measure	Interpretation
H1 Forecasting	Brier score, log loss and calibration at each horizon against common baselines	Ensemble must justify added complexity; report clustered uncertainty
H2 Compliance	Audited violation rate per observed vehicle-hour, before/after with controls	Separate behaviour change from increased detection coverage
H3 Financial viability	Lawful parking receipts and full operating costs; report penalty receipts separately	Test sustainability; no target to maximise penalties
H4 Search and emissions	Observed search duration and distance, with stated emission factors and counterfactual	Do not infer environmental savings from model accuracy
H5 Decision fairness	Incorrect-case rate, appeal outcomes and resolution time; audit by location and bay type	Interpret appeal rates with access and uptake, not as a standalone fairness score

Service targets remain separate from research findings: map-API latency at the 95th percentile of no more than 200 ms, sensor-to-map propagation within 10 seconds and 99.9% driver-service availability. These targets would be measured under declared loads and realistic failure conditions. Security, access controls, accessibility and retention would be tested as requirements; they are not certified properties of the present design.

4. RESULTS AND DISCUSSION

4.1 Design outputs and their interpretation

The output of this study is the connected specification in Figures 1-3 and Tables 1-4. It defines how a bay observation would be linked to a forecast, a service response and, where appropriate, an evidence record. There are no empirical prediction results or measured operational benefits to report. The proposed architecture is useful only to the extent that its interfaces and assumptions can be implemented and tested.

The shared bay record offers a way to trace disagreement between services. For example, if an application recommends a physically empty bay that is reserved for loading, the issue may lie in rule interpretation rather than occupancy prediction. If a bay is recommended from an old observation, the fault may instead be freshness handling. Separating these states helps assign failures to the component that needs correction and avoids reporting all unsuccessful recommendations as model errors.

This separation also limits the role of the ledger. Integrity verification can establish that a retained evidence package matches a previously recorded digest. It cannot validate a defective sensor or resolve an ambiguous restriction. A conventional signed append-only audit log should therefore be evaluated alongside a permissioned ledger. The comparison should cover recovery, access management, correction workflows, operating cost and the actual trust relationships between participating organisations.

4.2 Implications for kerbside operations

The system could support more targeted officer dispatch, but officers would still investigate exceptions, address safety problems and review disputed cases. Staffing benefits would need workload observations rather than assumptions about posts removed. The dashboard should show evidence quality and data gaps alongside possible contraventions so that a dense sensor installation does not automatically dominate operational attention.

Preparing machine-readable restrictions is itself a substantial task. Changes to loading hours, temporary suspensions or event rules must reach the operational model in time and retain their source. A nightly update can provide a routine reconciliation cycle but may be inadequate for urgent changes. The pilot should measure update delays and disable affected functions when authoritative status is uncertain.

For electric-vehicle bays, parking occupancy and charger operation are related but different. An empty space may have a failed charger, while a charging session may continue to appear active after its data feed becomes stale. Guidance would therefore need both bay and charger status, with separate timestamps. Accessible parking likewise requires information about eligibility and physical suitability; a high predicted vacancy alone is insufficient.

4.3 Fairness, privacy and incentives

A system used for public decisions must be understandable to people who do not use its application. Signage, payment alternatives, accessible communications and existing challenge routes remain necessary. The pilot should examine erroneous cases, time to resolution and user experience across locations and bay types, including areas with poor connectivity. Any demographic analysis would require a justified and lawful data collection plan rather than inferred personal attributes.

The governance problem extends beyond access control. More observations can increase the number of cases presented to officers even when compliance has improved. If success is defined by notice volume or income, the platform could reward that increase. The proposed evaluation instead separates compliance, incorrect decisions, lawful parking receipts and operating costs. These outcomes may move in different directions, and the study should report that rather than compress them into a single commercial score.

Sharing aggregated patterns between cities may reduce exposure of individual records, but aggregation alone does not establish anonymity. Transfer data would require a disclosure-risk assessment and controls on small groups, rare events and repeated releases. Similarly, opt-in app users may differ from other drivers. Findings from their journeys should be reported with the participation rate and the resulting limits on representativeness.

4.4 Limitations and reproducibility

The main limitation is the absence of empirical validation. Data availability, local permissions, label quality and implementation effort remain uncertain. Models developed in other cities provide methodological precedents rather than transferable performance guarantees. Seasonal demand, changes to tariffs, sensor failures and driver responses to guidance may all alter the problem after deployment. Guidance itself can redirect demand, so a model trained before launch may need reassessment once users act on its predictions.

The proposed ensemble adds maintenance work and may offer little advantage over a simpler baseline. Likewise, bay-level prediction may prove unreliable where vehicle turnover is high or observations are sparse. The pilot should permit a negative result: zone-level guidance, a simpler model or a conventional audit log may be more appropriate. Commercial viability and environmental benefit also remain open questions.

A future reproducibility package should include the bay schema, feature definitions, time cut-offs, model settings, split identifiers, training code, seeds and evaluation scripts. Data should be released only where rights and privacy safeguards permit. If raw records cannot be shared, documented synthetic examples and restricted-access procedures can support scrutiny, but they must be labelled clearly and must not be presented as the empirical pilot dataset. Protocol amendments and unsuccessful model variants should be recorded alongside favourable results.

5. CONCLUSION

ParkIQ proposes a shared bay-level architecture for parking guidance, occupancy forecasting, evidence review and operator services. The contribution is an engineering specification and an evaluation protocol that makes data provenance, uncertainty and decision accountability explicit. It does not establish prediction accuracy, financial returns or reductions in congestion and emissions.

The next step is a bounded pilot with checked bay records, authorised data access and independently audited observations. Simpler forecasting models and conventional audit logs should remain credible alternatives throughout that evaluation. Evidence from these comparisons would determine whether the proposed integration improves parking decisions and whether it is practical to maintain beyond a single deployment.

Declarations

Funding: No funding was received for this work.

Competing interests: The author declares no competing interests.

Data availability: No empirical dataset was generated or analysed for this design-and-evaluation protocol.

Human participants: This paper reports no human-participant study. Any subsequent pilot involving personal data or participants would require the relevant institutional and operational review before collection begins.

REFERENCES

- 1) Al Amiri, W., Baza, M., Banawan, K., Mahmoud, M., Alasmay, W., & Akkaya, K. (2019). Privacy-Preserving Smart Parking System Using Blockchain and Private Information Retrieval. 2019 International Conference on Smart Applications, Communications and Networking (SmartNets), 1-6. <https://doi.org/10.1109/smartnets48225.2019.9069783>
- 2) Barriga, J. J., Sulca, J., León, J. L., Ulloa, A., Portero, D., Andrade, R., & Yoo, S. G. (2019). Smart Parking: A Literature Review from the Technological Perspective. Applied Sciences, 9(21), 4569. <https://doi.org/10.3390/app9214569>

- 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) Cao, R., Choudhury, F., Winter, S., & Wang, D. Z. W. (2024). Artificial intelligence for parking forecasting: an extensive survey of machine learning techniques. *Transportmetrica A: Transport Science*, 1-39. <https://doi.org/10.1080/23249935.2024.2409229>
- 5) Channamallu, S. S., Kermanshachi, S., Rosenberger, J. M., & Pamidimukkala, A. (2026). Smart parking systems: A comprehensive review of digitalization of parking services. *Green Energy and Intelligent Transportation*, 5(1), 100293. <https://doi.org/10.1016/j.geits.2025.100293>
- 6) Department for Transport. (2022). Statutory guidance for local authorities in England on civil enforcement of parking contraventions. <https://www.gov.uk/government/publications/civil-enforcement-of-parking-contraventions/guidance-for-local-authorities-on-enforcing-parking-restrictions>
- 7) Dia, I., Ahvar, E., & Lee, G. M. (2022). Performance Evaluation of Machine Learning and Neural Network-Based Algorithms for Predicting Segment Availability in AIoT-Based Smart Parking. *Network*, 2(2), 225-238. <https://doi.org/10.3390/network2020015>
- 8) Enríquez, F. J., Mejía-Muñoz, J. M., Bravo, G., & Cruz-Mejía, O. (2024). Smart Parking: Enhancing Urban Mobility with Fog Computing and Machine Learning-Based Parking Occupancy Prediction. *Applied System Innovation*, 7(3), 52. <https://doi.org/10.3390/asi7030052>
- 9) Fahim, A., Hasan, M., & Chowdhury, M. A. (2021). Smart parking systems: comprehensive review based on various aspects. *Heliyon*, 7(5), e07050. <https://doi.org/10.1016/j.heliyon.2021.e07050>
- 10) Feng, Y., Xu, Y., Hu, Q., Krishnamoorthy, S., & Tang, Z. (2022). Predicting vacant parking space availability zone-wisely: a hybrid deep learning approach. *Complex & Intelligent Systems*, 8(5), 4145-4161. <https://doi.org/10.1007/s40747-022-00700-1>
- 11) Gama, J., Žliobaitė, I., Bifet, A., Pechenizkiy, M., & Bouchachia, A. (2014). A survey on concept drift adaptation. *ACM Computing Surveys*, 46(4), 1-37. <https://doi.org/10.1145/2523813>
- 12) Gneiting, T., Balabdaoui, F., & Raftery, A. E. (2007). Probabilistic Forecasts, Calibration and Sharpness. *Journal of the Royal Statistical Society Series B: Statistical Methodology*, 69(2), 243-268. <https://doi.org/10.1111/j.1467-9868.2007.00587.x>
- 13) Hao, H., Wang, Y., Du, L., & Chen, S. (2023). Enabling smart curb management with spatiotemporal deep learning. *Computers, Environment and Urban Systems*, 99, 101914. <https://doi.org/10.1016/j.compenvurbsys.2022.101914>
- 14) Information Commissioner's Office. (n.d.). Automatic Number Plate Recognition (ANPR). Retrieved September 24, 2026, from <https://ico.org.uk/for-organisations/uk-gdpr-guidance-and-resources/cctv-and-video-surveillance/guidance-on-video-surveillance-including-cctv/additional-considerations-for-technologies-other-than-cctv/automatic-number-plate-recognition-anpr/>
- 15) INRIX. (2017, July 12). Searching for parking costs the UK £23.3 billion a year. <https://inrix.com/press-releases/parking-pain-uk/>
- 16) Kapoor, S., & Narayanan, A. (2023). Leakage and the reproducibility crisis in machine-learning-based science. *Patterns*, 4(9), 100804. <https://doi.org/10.1016/j.patter.2023.100804>
- 17) Lin, T., Rivano, H., & Le Mouel, F. (2017). A Survey of Smart Parking Solutions. *IEEE Transactions on Intelligent Transportation Systems*, 18(12), 3229-3253. <https://doi.org/10.1109/tits.2017.2685143>
- 18) Liu, Y., Liu, C., & Luo, X. (2021). Spatiotemporal Deep-Learning Networks for Shared-Parking Demand Prediction. *Journal of Transportation Engineering, Part A: Systems*, 147(6), 04021026. <https://doi.org/10.1061/jtepbs.0000522>
- 19) Ma, C., Huang, X., & Li, J. (2024). A review of research on urban parking prediction. *Journal of Traffic and Transportation Engineering (English Edition)*, 11(4), 700-720. <https://doi.org/10.1016/j.jtte.2023.11.004>
- 20) Navpreet, Rani, R., & Roul, R. K. (2025). State-of-the-Art Machine Learning and Deep Learning Techniques for Parking Space Classification: A Systematic Review. *Archives of Computational Methods in Engineering*, 32(6), 3851-3883. <https://doi.org/10.1007/s11831-025-10250-7>
- 21) Perković, T., Šolić, P., Zargariasl, H., Čoko, D., & Rodrigues, J. J. P. C. (2020). Smart Parking Sensors: State of the Art and Performance Evaluation. *Journal of Cleaner Production*, 262, 121181. <https://doi.org/10.1016/j.jclepro.2020.121181>
- 22) RAC Foundation. (2025, November 19). Council parking revenue in England 2024-25. <https://www.racfoundation.org/research/economy/council-parking-revenue-in-england-2024-25>
- 23) Roberts, D. R., Bahn, V., Ciuti, S., Boyce, M. S., Elith, J., Guillera-Aroita, G., Hauenstein, S., Lahoz-Monfort, J. J., Schröder, B., Thuiller, W., Warton, D. I., Wintle, B. A., Hartig, F., & Dormann, C. F. (2017). Cross-validation strategies for data with temporal, spatial, hierarchical, or phylogenetic structure. *Ecography*, 40(8), 913-929. <https://doi.org/10.1111/ecog.02881>
- 24) Rocco, G., Pipino, C., & Pagano, C. (2023). An Overview of Urban Mobility: Revolutionizing with Innovative Smart Parking Systems. *Sustainability*, 15(17), 13174. <https://doi.org/10.3390/su151713174>
- 25) Sha, H., Haouari, R., Singh, M. K., Papazikou, E., Quddus, M., Chaudhry, A., Thomas, P., & Morris, A. (2024). How can on-street parking regulations affect traffic, safety, and the environment in a cooperative, connected, and automated era?. *European Transport Research Review*, 16(1), 18. <https://doi.org/10.1186/s12544-023-00628-8>
- 26) Shoup, D. C. (2006). Cruising for parking. *Transport Policy*, 13(6), 479-486. <https://doi.org/10.1016/j.tranpol.2006.05.005>
- 27) Vlahogianni, E. I., Kepaptsoglou, K., Tsetsos, V., & Karlaftis, M. G. (2016). A Real-Time Parking Prediction System for Smart Cities. *Journal of Intelligent Transportation Systems*, 20(2), 192-204. <https://doi.org/10.1080/15472450.2015.1037955>
- 28) Xiao, J., Lou, Y., & Frisby, J. (2018). How likely am I to find parking? - A practical model-based framework for predicting parking availability. *Transportation Research Part B: Methodological*, 112, 19-39. <https://doi.org/10.1016/j.trb.2018.04.001>
- 29) Xiao, X., Peng, Z., Lin, Y., Jin, Z., Shao, W., Chen, R., Cheng, N., & Mao, G. (2023). Parking Prediction in Smart Cities: A Survey. *IEEE Transactions on Intelligent Transportation Systems*, 24(10), 10302-10326. <https://doi.org/10.1109/tits.2023.3279024>

- 30) Yang, H. F., Ke, R., Cui, Z., Wang, Y., & Murthy, K. (2022). Toward a real-time Smart Parking Data Management and Prediction (SPDMP) system by attributes representation learning. *International Journal of Intelligent Systems*, 37(8), 4437-4470. <https://doi.org/10.1002/int.22725>
- 31) Yang, S., Ma, W., Pi, X., & Qian, S. (2019). A deep learning approach to real-time parking occupancy prediction in transportation networks incorporating multiple spatio-temporal data sources. *Transportation Research Part C: Emerging Technologies*, 107, 248-265. <https://doi.org/10.1016/j.trc.2019.08.010>
- 32) Zeng, F., Pang, C., & Tang, H. (2024). Sensors on Internet of Things Systems for the Sustainable Development of Smart Cities: A Systematic Literature Review. *Sensors*, 24(7), 2074. <https://doi.org/10.3390/s24072074>
- 33) Zulfiqar, H., Ul Haque, H. M., Tariq, F., & Khan, R. M. (2023). A survey on smart parking systems in urban cities. *Concurrency and Computation: Practice and Experience*, 35(15), e6511. <https://doi.org/10.1002/cpe.6511>