Journal of Management and Planning Research
2024, 1(1), 190-199

DATA-DRIVEN PARKING SEARCH INTELLIGENCE FOR URBAN MANAGEMENT AND PLANNING: EVIDENCE FROM LARGE-SCALE VEHICLE GPS TRACES IN FRANKFURT

Andrew Saint
Naveed Ahmad

Cruising for parking remains a localized but consequential management problem because it adds avoidable delay, concentrates congestion around high-demand destinations, and weakens the operational performance of curbside systems. This paper presents a management- and planning-oriented synthesis of a validated GPS-based parking-search identification framework and evaluates what that established measurement approach contributes to operational decision-making using three linked empirical components: a labeled ground-truth corpus of 3,550 journeys, an external dynamic park-and-visit dataset of 161 journeys, and a filtered large-scale archive of 868,561 consumer light-vehicle trips ending in Frankfurt. The results show that the predictive model achieves a mean absolute error below one minute across sampling rates from 1 to 15 seconds, reduces prediction error materially relative to a constant-duration baseline, and remains robust when transferred to an external dataset collected with a different application. When applied at city scale, the model indicates that 33% of trips end with immediate parking, while the overall mean parking-search duration is 1 minute 30 seconds and rises to 2 minutes 15 seconds among trips with non-zero search. The evidence also shows a substantial reduction in vehicle speed near parking-search activity and identifies central districts with systematically elevated mean search duration. Taken together, the results show that GPS-derived parking-search intelligence can be translated into credible, decision-ready evidence for targeted curb regulation, parking guidance, infrastructure prioritization, and more realistic travel-time planning.

---

https://doi.org/10.66033/jmpr2024-117

INTRODUCTION

Cruising for parking is not merely a nuisance at the end of a trip. It is a recurring urban management problem that intensifies congestion in already sensitive locations, increases emissions at the street level, and degrades the reliability of urban access systems [10, 5, 12]. For planners and managers, the core challenge is measurement. If parking-search behavior cannot be identified credibly, policy design is forced to rely on self-reports, narrow field experiments, or heuristic assumptions that are difficult to scale and difficult to trust.

Recent work on GPS-based parking-search identification substantially improves this situation by introducing directly labeled ground-truth data and a predictive model capable of distinguishing normal driving from parking-search behavior in vehicle trajectories [3]. For management and planning research, the key importance of that advance is not novelty for its own sake, but the fact that it establishes a credible measurement pipeline linking behavioral detection to operational decision-making. Once parking-search duration can be measured at scale, it becomes possible to identify where parking supply is functionally misaligned with demand, when curb pressure intensifies, and where guidance or pricing interventions are likely to be most effective [4].

The present manuscript is organized around that planning problem. Its contribution is therefore applied and translational: rather than claiming a wholly new detection architecture, it clarifies the managerial significance, empirical reliability, and planning usefulness of an already validated identification framework. The paper addresses three questions:

RQ1: How reliable is the validated GPS-based parking-search model across sampling conditions and data sources?

RQ2: What do the reported empirical results reveal about the scale and distribution of parking-search burden in Frankfurt?

RQ3: How can these results be translated into actionable implications for parking management and urban planning?

The topic is well aligned with management and planning scholarship because it connects a robust measurement framework with questions of resource allocation, operational design, policy timing, and infrastructure governance. The contribution of the present manuscript lies in that decision-oriented framing: it translates a validated technical model into a management instrument that can inform planning analysis and operational policy.

LITERATURE AND ANALYTICAL ORIENTATION

The study of parking search has historically been constrained by two weaknesses. First, survey-based estimates are often affected by recall bias and limited representativeness [1, 3, 2]. Second, many trajectory-based approaches identify parking search through simple assumptions, such as entering a fixed radius around the parking location or crossing a speed threshold. These methods are easy to implement, but they often fail to capture the behavioral complexity of real-world parking decisions [5].

The development of a data-driven identification model trained on directly labeled trips changes the research frontier in two important ways [6]. Methodologically, it reduces dependence on arbitrary heuristics. Substantively, it shifts attention toward management use cases: parking search can now be analyzed as an operational indicator of curb inefficiency, access friction, and infrastructure mismatch. In practical terms, this matters for at least four reasons.

- It improves the credibility of aggregate statistics used in planning reports and policy design.
- It enables comparisons between search conditions across times, locations, and data sources.
- It provides an empirical basis for guidance systems, curb regulation, and parking supply management.
- It makes generalized, large-scale monitoring feasible without requiring constant field observation.

This analytical orientation is especially relevant to management and planning because the value of the model is not exhausted by predictive accuracy alone. The model itself is already technically validated; the contribution here is to assess how reliably that evidence can be interpreted for planning use and to show how trajectory data can be converted into decision-relevant operational intelligence.

DATA SOURCES AND METHODOLOGICAL FRAMEWORK

Ground-truth corpus

The empirical foundation begins with a purpose-built smartphone application designed to record four critical journey phases with direct user input: trip origin, the start of parking search, the parking spot, and the final destination. After cleaning more than 7,000 initiated journeys, the valid ground-truth corpus comprises 3,550 trips recorded by 162 drivers in Germany. The dataset is geographically anchored in Frankfurt, with 2,344 journeys (about 66%) ending there [7].

The descriptive properties of this corpus are important for both model training and managerial interpretation. Approximately 18% of journeys recorded a parking-search duration (PSD) of zero, indicating immediate parking. Across all journeys, the mean PSD is 1 minute 25 seconds; among trips with PSD greater than zero, the mean rises to 1 minute 44 seconds [8]. The mean walking duration is 2 minutes 40 seconds across all trips and 2 minutes 52 seconds among trips with positive PSD. The 95th percentile of PSD across all journeys is 4 minutes 59 seconds, showing that long searches occur but are not typical. For trips with positive PSD, the mean Initial Search Radius (ISR) is 137 meters, while the mean Parking Offset Radius (POR) is 128 meters. The underlying GPS records include roughly 2.29 million points, mostly at a 1-second sampling rate [9].

Predictive model structure

The predictive model is a feed-forward neural network built to identify the onset and continuation of parking search within a trajectory. Its architecture contains two hidden layers with 128 and 32 neurons, followed by a sigmoid output that yields the probability that a given GPS point belongs to the parking-search phase [10]. For each validation run, the classification cutoff is selected from the training portion of the data to minimize mean absolute error in predicted PSD and is then applied unchanged to the held-out trips [11].

A key strength of the model is parsimony. After feature testing, the final model relies on only three substantive variables: distance to parking spot, vehicle speed, and sampling rate. These are expanded into 13 input features through lag structure: one contemporaneous distance measure, six speed measures (current plus five lags), and six sampling-rate measures (current plus five lags). This design improves portability because it does not depend on context-specific variables that are often absent in historical GPS archives, while also reducing the risk that apparent performance is driven by highly localized covariates that are unavailable in large operational archives [12].

Table 1: Key descriptive statistics from the labeled ground-truth corpus

Measure	Count	Reported value
Valid journeys after cleaning	3,550	More than 7,000 initiated trips were reduced to a cleaned analytic corpus.
Drivers	162	Average of approximately 22 valid journeys per driver.
Journeys ending in Frankfurt	2,344	About 66% of the corpus.
Trips with immediate parking (PSD ≈ 0)	626	About 18% of all journeys.
Mean PSD (all journeys)	3,550	01:25
Mean PSD (journeys with PSD > 0)	2,924	01:44
Median PSD (all journeys)	3,550	00:50
95th percentile PSD (all journeys)	3,550	04:59
Mean walking duration (all journeys)	3,550	02:40
Mean walking duration (journeys with PSD > 0)	2,924	02:52
Mean Initial Search Radius (ISR)	2,924	137 m
Mean Parking Offset Radius (POR)	2,924	128 m
GPS points in corpus	2,293,850	Mostly recorded at a 1-second sampling rate.
Average speed within 1 km of destination (normal driving)	492,355 points	30 km/h
Average speed within 1 km of destination (parking search)	123,482 points	16 km/h

Note: All values are reported from the labeled training corpus. Duration values are reported in minutes:seconds.

Algorithm 1 Operational workflow for GPS-based parking-search identification

Require: GPS trajectory points with timestamps, speed values, and parking spot location
Ensure: Predicted parking-search phase and parking-search duration for each trip

```
for each trip trajectory do
    Compute the current distance from each GPS point to the final parking spot
    Construct the 13-feature input vector using distance, speed, and sampling-rate history
    Pass each point through the trained feed-forward neural network
    Obtain a parking-search probability for every GPS point
    Apply the sampling-rate-specific probability cutoff selected during validation
    Identify the first transition from normal driving to parking search
    Label all subsequent points as parking-search points
    Convert the labeled segment into predicted parking-search duration
end for
Aggregate trip-level outputs for validation, benchmarking, and planning analysis
```

Validation design

Validation was conducted in three complementary ways. First, the model was assessed under 10-fold cross-validation using resampled data at 1-, 5-, 10-, and 15-second intervals. Because the labeled corpus contains repeated trips by the same participants, fold assignment was specified at the driver level so that all journeys from a given driver remained within a single fold, thereby reducing the risk of dependence leakage between training and validation data. Second, performance was compared against a constant-duration baseline that always predicts the overall mean PSD. Third, the model was tested on an external dynamic park-and-visit dataset collected with a different application and different drivers.

The external dataset contains 161 journeys recorded between November 2020 and June 2021, primarily in Frankfurt (96 journeys) and Rostock (63 journeys). The protocol intentionally allowed drivers to choose when to begin searching and where to park, making the test closer to real driving behavior than fixed-start experiments. This external comparison is important because it supplements the cross-validation evidence with a genuinely out-of-sample test collected under a different recording process. Within this dataset, 57 trips (about 35%) had a PSD of zero; the mean PSD was 1 minute 20 seconds, the median was 30 seconds, the 75th percentile was 1 minute 40 seconds, and the maximum observed PSD was 22 minutes.

EMPIRICAL RESULTS

Validation and model reliability

Across sampling conditions, the model maintains a mean absolute error of less than one minute and shows only negligible deviation between actual and predicted mean PSD. This is important for planning practice because large operational datasets rarely share a uniform sampling frequency. The reported performance indicates that the model remains usable even when data are less granular than the original 1-second corpus. No single metric is sufficient on its own, so these results are interpreted jointly with agreement in mean PSD, comparison against heuristic baselines, and external testing rather than as a stand-alone claim of complete model adequacy.

Table 2: Model validation across sampling rates and external testing

Validation setting	**MAE**	**MPSD (actual)**	**MPSD (pred.)**	**Interpretation**
1-second sampling	00:52	01:26	01:27	Highest granularity; near-exact aggregate fit.
5-second sampling	00:52	01:26	01:27	No practical degradation relative to 1-second data.
10-second sampling	00:54	01:26	01:26	Stable aggregate accuracy with modestly higher individual error.
15-second sampling	00:56	01:26	01:23	Slightly lower precision, but still below one minute MAE.
Constant-duration baseline (1-second comparison)	01:20	01:26	01:26	Predicting the same duration for every trip performs materially worse than the neural model.
Trips within Frankfurt	≈00:52	—	—	Error remains low in the denser urban core.
Trips outside Frankfurt	≈00:40	—	—	Slightly lower error suggests geographic portability.
External dynamic park-and-visit dataset	00:39	—	—	Mean PSD deviation between actual and predicted values is only 17 seconds.

Note: Duration values are reported in minutes:seconds. The constant-duration baseline always predicts the overall mean PSD of 01:26 for every observation.

The baseline comparison is especially informative from a management perspective. A simple average-based rule yields an MAE of approximately 80 seconds, whereas the neural model reduces this to 52 seconds in the 1-second validation setting. That 28-second improvement per trip is operationally meaningful when aggregated across thousands of daily arrivals, particularly because it improves estimation precisely in the terminal segment of trips where curbside pressure is most concentrated.

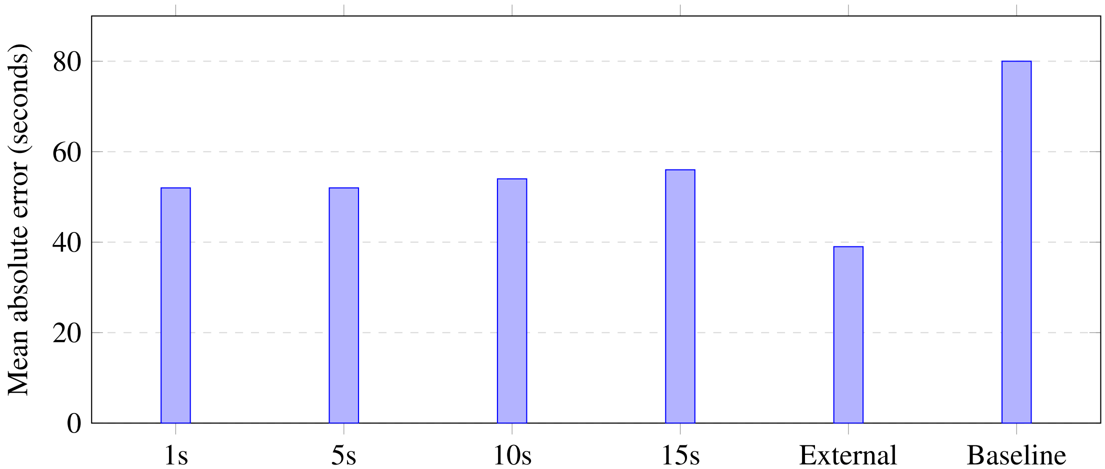

Figure 1: Reported mean absolute error across validation settings. The external test uses the dynamic park-and-visit dataset; the baseline always predicts a constant duration.

Benchmarking against heuristic alternatives

The model was also benchmarked against four alternative identification approaches commonly used in the literature: a naive 200-meter radius rule, a speed-threshold rule, a first-local-minima rule, and an actual-shortest-path comparison. This benchmarking step serves as an additional validation layer because it shows that the model is not only accurate in isolation, but also materially more reliable than commonly used heuristic rules on both trip-level error and aggregate duration estimates.

Table 3: Comparative performance of parking-search identification methods

Method	MAE	Mean PSD	Median PSD	Mean POR (m)
Ground truth	—	01:25	00:50	120
Naive 200 m radius	01:06	02:21	01:40	—
Speed threshold	01:22	02:38	01:56	225
First local minima	01:10	01:39	00:47	68
Actual-shortest path	—	02:05	01:46	—
Machine-learning model	00:48	01:26	01:08	87

Note: Duration values are reported in minutes:seconds. The actual-shortest-path method estimates excess travel time and is therefore not directly comparable on MAE to the directly labeled ground truth.

Two findings stand out. First, the machine-learning model achieves the lowest directly comparable MAE at 48 seconds. Second, the aggregate mean PSD generated by the model (1 minute 26 seconds) is almost identical to the ground-truth mean (1 minute 25 seconds). This combination of individual-level and aggregate-level accuracy is what makes the model useful for planning applications. It is not simply a higher-performing technical classifier in this comparison; it is a more credible operational measurement instrument for estimating terminal parking friction.

Large-scale evidence from Frankfurt

The large-scale planning value of the framework emerges most clearly in the Frankfurt application. The original INRIX archive contained approximately 18 million journeys with billions of GPS points associated with a bounding box around Frankfurt in 2019. After filtering to consumer trips, light-weight vehicles, and trips ending in Frankfurt, the analytic dataset contains 868,561 journeys.

The filtered archive is substantial not only in size but also in heterogeneity. The median sampling rate is 6 seconds, the mean journey distance is 40 km, the median journey distance is 15 km, the mean journey duration is 35 minutes, and mean journey speed is 49 km/h. These characteristics make the archive well suited to city-scale planning analysis.

Table 4: Reported large-scale descriptive evidence from the Frankfurt application

Measure	**Count**	**Reported value**
Filtered trips ending in Frankfurt	868,561	Consumer trips, light-weight vehicles, 2019.
Trips with immediate parking (PSD = 0)	285,594	33% of all trips.
Trips with positive PSD	582,967	67% of all trips.
Mean PSD (all trips)	868,561	01:30
Median PSD (all trips)	868,561	00:15
95th percentile PSD (all trips)	868,561	08:20
Mean PSD (trips with PSD > 0)	582,967	02:15
Median PSD (trips with PSD > 0)	582,967	00:42
95th percentile PSD (trips with PSD > 0)	582,967	10:30
Mean Parking Offset Radius (positive PSD)	582,967	143 m
Median Parking Offset Radius (positive PSD)	582,967	127 m
95th percentile Parking Offset Radius (positive PSD)	582,967	309 m
Average speed within 1 km of parking spot (normal driving)	23,324,499 points	31 km/h
Average speed within 1 km of parking spot (parking search)	8,832,360 points	17 km/h

Note: Duration values are reported in minutes:seconds. The large-scale archive was filtered from an initial INRIX dataset of approximately 18 million journeys.

Three substantive planning findings follow from these reported results.

First, immediate parking is common but far from universal. One-third of trips end without observable search, which means two-thirds still experience some degree of terminal friction. This matters for access planning because average citywide conditions can appear manageable while still masking substantial pressure during many arrivals.

Second, the PSD distribution is strongly right-skewed. The median trip ends with only 15 seconds of search, yet the 95th percentile reaches 8 minutes 20 seconds for the full sample. For managers, this means the average is only part of the story. The upper tail captures the operational stress most visible to drivers, retailers, and neighborhood streets.

Third, parking search is associated with a marked behavioral shift in vehicle speed. Within one kilometer of the parking location, average speed falls from 31 km/h during normal driving to 17 km/h during parking search. This confirms that parking search is not simply extra time; it is a distinct, slow-moving operational state that can intensify local traffic frictions.

The spatial analysis reported in the Frankfurt application identifies particularly high mean parking-search duration in the *Innenstadt* and *Altstadt*. This is notable because those areas already contain multiple parking facilities, suggesting that the problem is not purely one of nominal supply. At the same time, the archive does not directly observe occupancy, pricing, visibility, or driver preference, so these mechanisms should be interpreted as plausible explanations for the observed pattern rather than identified causal drivers.

MANAGEMENT AND PLANNING IMPLICATIONS

The empirical evidence supports several management and planning conclusions.

Curb management should be targeted, not generic

Because the upper tail of PSD is substantially larger than the median, managers should avoid relying solely on citywide averages. Policies are likely to be more effective when they focus on the locations and time windows where the operational burden is most visible, particularly high-demand central districts and destination-rich commercial zones.

Parking guidance has direct operational value

The combination of near-real-time detectability and strong aggregate accuracy makes the framework useful for parking guidance systems. If navigation platforms can identify or anticipate parking search near the end of a trip, they can redirect drivers toward available structured parking before slow search loops begin. The reported evidence therefore supports, rather than by itself proves, the operational value of information provision in reducing curbside cruising [6, 7].

Travel-time estimation should include terminal parking friction

Most route planning tools treat arrival at the destination street as the end of the trip. The reported evidence shows why this is incomplete. Even in a citywide average sense, drivers spend 1 minute 30 seconds in additional terminal search time, and non-zero search episodes average 2 minutes 15 seconds. Incorporating this friction into generalized travel-time estimates would improve trip planning, mode comparison, and accessibility assessment.

Off-street supply must be made legible and behaviorally reachable

The persistence of high search duration in central districts despite the presence of garages indicates that supply alone does not necessarily resolve parking search. For planning practice, this points toward improved wayfinding, clearer pricing signals, dynamic signage, and better last-block routing to structured parking. Management failure can arise even when infrastructure exists but is not effectively legible or behaviorally reachable.

LIMITATIONS

Several limitations remain important for interpretation. First, the ground-truth corpus is based on active volunteer participation and cannot be treated as a representative sample of all drivers. Second, the data collection process can influence behavior, including the possibility that participants alter their parking conduct because they know they are being observed. Third, in historical GPS archives the exact pedestrian destination is often unavailable, so Parking Offset Radius functions only as a proxy for walking tolerance. Fourth, the large-scale results provide strong descriptive evidence but do not directly observe pricing, occupancy constraints, legal restrictions, or individual preferences. Fifth, the validation evidence is centered on mean absolute error, agreement in mean PSD, benchmarking, and external transfer; these are appropriate and informative diagnostics for trip-level operational measurement, but they do not substitute for every possible assessment of calibration, threshold sensitivity, or uncertainty. These limitations do not invalidate the framework, but they do define the boundary between robust operational measurement and fuller causal explanation.

CONCLUSION

This paper demonstrates the management and planning relevance of validated GPS-based parking-search identification. The evidence shows that a parsimonious neural model built from directly labeled trips can identify parking-search behavior with sub-minute mean absolute error across multiple sampling rates, outperform simpler heuristic approaches, and remain informative in an external validation dataset collected under a different protocol. When deployed at scale in Frankfurt, the model reveals that one-third of trips end with immediate parking, while the remaining trips generate a meaningful layer of terminal delay, with an overall mean parking-search duration of 1 minute 30 seconds and substantially longer delays in the upper tail.

For management and planning research, the main contribution is practical and interpretive. The manuscript does not depend on claiming a wholly new machine-learning architecture; instead, it shows why an already validated identification framework is consequential for operational planning. Parking search can now be measured in a way that is sufficiently robust for city-scale use, making it possible to design better curb policies, improve garage guidance, incorporate terminal parking friction into accessibility analysis, and prioritize interventions where local search burden is highest. In that sense, the importance of the framework lies not only in prediction, but in its ability to convert ubiquitous mobility traces into decision-ready evidence for urban governance.

REFERENCES

[1] Alemi, F., Rodier, C., and Drake, C. (2018). Cruising and on-street parking pricing: A difference-in-difference analysis of measured parking search time and distance in San Francisco. *Transportation Research Part A: Policy and Practice*, 111, 187–198.

[2] Assemi, B., Baker, D., and Paz, A. (2020). Searching for on-street parking: An empirical investigation of the factors influencing cruising time. *Transport Policy*, 97, 186–196.

[3] Belloche, S. (2015). On-street parking search time modelling and validation with survey-based data. *Transportation Research Procedia*, 6, 313–324.

[4] Bisante, A., Panizzi, E., and Zeppieri, S. (2023). Cruising-for-parking detection on the smartphone based on implicit interaction and machine learning. In *Proceedings of the 15th International Conference on Automotive User Interfaces and Interactive Vehicular Applications*, 93–102.

[5] Brooke, S., Ison, S., and Quddus, M. (2014). On-street parking search. *Transportation Research Record*, 2469(1), 65–75.

[6] Dalla Chiara, G., Krutein, K. F., Ranjbari, A., and Goodchild, A. (2022). Providing curb availability information to delivery drivers reduces cruising for parking. *Scientific Reports*, 12, 19355.

[7] Fahim, A., Hasan, M., and Chowdhury, M. A. (2021). Smart parking systems: Comprehensive review based on various aspects. *Heliyon*, 7(5), e07050.

[8] Millard-Ball, A., Hampshire, R. C., and Weinberger, R. (2020). Parking behaviour: The curious lack of cruising for parking in San Francisco. *Land Use Policy*, 91, 103918.

[9] Saki, S., and Hagen, T. (2024). Parking search identification in vehicle GPS traces. *Journal of Urban Mobility*, 6, 100083.

[10] Shoup, D. C. (2005). *The High Cost of Free Parking*. Planners Press.

[11] Shoup, D. (2021). Pricing curb parking. *Transportation Research Part A: Policy and Practice*, 154, 399–412.

[12] Zhu, Y., Ye, X., Chen, J., Yan, X., and Wang, T. (2020). Impact of cruising for parking on travel time of traffic flow. *Sustainability*, 12(8), 3079.

Andrew Saint, Department of Architecture, University of Cambridge

Naveed Ahmad, Norwegian University of Life Sciences

Manuscript Published; 10 December 2024.