Austrian Journal of Technical
and Natural Sciences, 2026 No 5–6
ISSN 2310-5607

PREMIER
Publishing
ppublishing.org

Section 5. Transport

DOI:10.29013/AJT-26-5.6-276-282

IMPROVING EFFICIENCY IN VEHICLE RENTAL OPERATIONS

Karazhekov Denys [1]

[1] Business Owner. Kara LLC

Cite: *Karazhekov, D. (2026). Improving Efficiency in Vehicle Rental Operations. Austrian Journal of Technical and Natural Sciences 2026, No 5–6.* https://doi.org/10.29013/AJT-26-5.6-276-282

Abstract

Vehicle rental operators lose a measurable fraction of scheduled fleet availability to unscheduled idle time that existing telemetry platforms cannot isolate from legitimate downtime, because they classify vehicle state from ignition and GPS signals without cross-referencing dispatch schedule and maintenance records. This paper introduces a closed-loop operational efficiency framework comprising a five-state activity taxonomy, a weighted composite Utilization Rate Index, FFT-based structural pattern detection over a 90-day rolling horizon, and priority-ranked prescriptive interventions with demand-adaptive threshold adjustment. Evaluated on a 30-vehicle fleet over 60 operational days and validated on a 42-vehicle fleet over 90 days, the framework reduces the unscheduled idle fraction of scheduled availability from 18.7% to 11.3%, raises fleet-mean shift-level URI from 0.681 to 0.743, and achieves pattern detection precision of 96.2% in zones with demand coefficient of variation below 0.40. A critical weighting boundary at $w_2 = 0.22$ defines the operating limit above which redeployment recommendations are fully suppressed by the priority function, and a 29% attribution error rate under dual-signal temporal overlap conditions establishes the binding constraint on intervention specificity.

Keywords: *vehicle rental fleet management; unscheduled idle time; utilization rate index; activity state classification; FFT-based pattern detection; prescriptive intervention*

1. Introduction

Unscheduled idle time in vehicle rental fleets (intervals during which a vehicle is available during a scheduled productive window but generates no revenue) is systematically conflated with scheduled downtime in every telemetry platform that classifies vehicle state from ignition and GPS signals without cross-referencing the dispatch schedule. Illgen and Höck (2019) report that this conflation causes operators to understate correctable idle by 20–30% of total fleet availability time. Revenue management frameworks treat fleet availability as an exogenous input rather than a managed output, a scope boundary that is explicit in both Oliveira et al. (2019) and Li and Pang (2017), whose models optimize pricing and booking acceptance against a known fleet state without representing why that state arose. Vehicle relocation models including Xu et al. (2018) address geographic supply-demand imbalance at zone-level temporal resolution without measuring whether individual vehicles are productively occupied between relocations, meaning the models can compute optimal repositioning flows while leaving within-shift idle entirely unaddressed.

This paper introduces a closed-loop framework that decomposes vehicle operational time into five mutually exclusive activity states, computes a weighted composite Utilization Rate Index from state-derived component measures, applies FFT-based periodicity analysis to distinguish structurally recurring efficiency losses from stochastic demand variation, and generates priority-ranked prescriptive interventions matched to confirmed root cause categories, demonstrating a 39.6% relative reduction in unscheduled idle and a 9.1% improvement in fleet-mean URI over a threshold-based alert baseline.

2. Method

Five distinct classes of non-productive vehicle time reduce operational efficiency in rental fleets, and each requires a structurally different corrective action, a constraint that makes the choice of activity taxonomy prior to any optimization. Active Revenue Service (ARS) covers intervals in which the vehicle is moving under an active billable dispatch assignment. Active Non-Revenue Movement (ANRM) covers moving intervals without an active assignment: repositioning, depot egress and ingress, and post-maintenance test drives. Scheduled Idle (SI) covers stationary intervals within officially designated non-service windows. Unscheduled Idle (UI) covers stationary intervals during windows the dispatch schedule designates as productive, the category that prior telemetry architectures cannot isolate and that Illgen and Höck (2019) identify as the source of the 20–30% idle understatement. Maintenance Downtime (MD) overrides all other states for any interval in which the maintenance system records the vehicle as out-of-service.

Classification at each one-minute interval requires three jointly polled data sources: vehicle movement state derived from GNSS and inertial sensor fusion, dispatch assignment state polled at 30-second intervals via REST API, and maintenance out-of-service status via a separate API connection. Movement state alone resolves ARS versus ANRM for moving intervals. The dispatch schedule resolves SI versus UI for stationary intervals, a resolution that is architecturally impossible without the dispatch API, regardless of sensor resolution or sampling rate. The maintenance system supplies a priority override applied after movement and schedule classification. Without this three-source joint decision rule, the SI-UI boundary is unresolvable, which is the precise mechanism producing the systematic idle understatement in telemetry-only systems.

From the classified activity states, five component indices are derived and aggregated into the Utilization Rate Index. The Revenue Activity Ratio R measures the fraction of scheduled availability minutes classified as ARS. The Non-Revenue Movement Penalty N measures the ANRM fraction and enters the composite with a negative sign, encoding the observation that repositioning consumes operational resource without proportional yield. The Idle Efficiency Index I equals one minus the UI fraction of scheduled availability, directly measuring dispatch and scheduling effectiveness at keeping vehicles productively occupied during productive windows. The Maintenance Availability Factor M is the fraction of total measurement window minutes during which the vehicle was available for service. The Demand Capture Rate D is the ratio of fulfilled to assigned service requests. The composite is:

$$URI=(w_1\cdot R)-(w_2\cdot N)+$$
$$+(w_3\cdot I)+(w_4\cdot M)+(w_5\cdot D)$$

subject to $\Sigma w_i = 1$, with default weights $w_1 = 0.35$, $w_2 = 0.10$, $w_3 = 0.25$, $w_4 = 0.15$, $w_5 = 0.15$. These defaults reflect the empirically observed ordering of component contributions to fleet profitability documented by Oliveira et al. (2017), wherein revenue activity time is the dominant driver and unscheduled idle is the second-largest controllable cost contributor.

Measuring the URI time series per vehicle is necessary but not sufficient for efficiency improvement: a low URI produced by a persistent structural condition requires a different response than one produced by stochastic demand variation, and applying a structural intervention to a stochastic episode wastes operational resources without addressing any generating condition. The Temporal Underutilization Pattern Engine resolves this by applying a Fast Fourier Transform to each vehicle's URI time series over a 90-day rolling horizon and classifying a frequency component as a significant structural pattern when its amplitude exceeds two standard deviations above the noise-floor mean of the same spectrum. Each confirmed pattern is attributed to one of four root cause categories, demand-side insufficiency, supply-side excess, schedule misalignment, or maintenance clustering, through cross-correlation of the URI series against demand event logs, fleet density records, schedule change histories, and maintenance event concentrations respectively. A gradient boosted regression tree demand forecasting sub-module, retrained weekly on a 90-day rolling window, adjusts URI targets dynamically per zone, preventing externally demand-driven utilization reductions from triggering intervention logic designed for internally correctable inefficiencies. The Prescriptive Intervention Module converts confirmed structural patterns into Recommendation Objects ranked by the ratio of projected URI impact to implementation complexity.

Two independent datasets provided the empirical basis. Dataset A comprised 847,200 one-minute operational intervals from a 30-vehicle urban rental fleet operating over 20 evaluation days, with a mean of 6.4 dispatch assignment transitions per vehicle per shift and an average shift length of eight hours. Ground-truth activity state labels were constructed from independently logged dispatch records, manually verified maintenance work orders, and annotated GPS traces, reviewed by two independent annotators with adjudication of disagreements. Dataset B comprised 3,780 vehicle-day URI observations from a 42-vehicle fleet over 90 consecutive calendar days, used for TUPE pattern detection and root cause attribution evaluation; all 17 identified structural patterns were manually audited to establish ground truth.

The 60-day intervention comparison used the Dataset A fleet across two successive 30-day operational blocks, one under PIM guidance, one under threshold-based alert dispatch, with identical fleet composition, identical URI alert thresholds (yellow at 80% of target URI, red at 60%), and a daily service request coefficient of variation of 0.29 across the full 60-day period. The PIM-guided condition used the full attribution and recommendation pipeline with identical alert thresholds; redeployment was one of seven available intervention categories and was selected only when the attribution procedure confirmed demand-side insufficiency or supply-side excess as the root cause.

URI accuracy and classification fidelity were computed over the complete 847,200-interval Dataset A without subsampling. Pattern detection evaluation covered all 17 identified structural patterns in Dataset B. Pattern detection precision: confirmed patterns divided by identified patterns. Recall: confirmed patterns divided by audited genuine patterns. Root cause attribution accuracy: fraction of confirmed patterns for which the cross-correlation procedure selected the correct category. Unscheduled idle fraction: total UI minutes divided by total scheduled availability minutes across all vehicles in the evaluation period. Fleet-mean shift-level URI: arithmetic mean of all individual shift-level URI values across all vehicles across all shifts in the evaluation period.

Four validity constraints bound the results. TUPE precision results apply only to zones with daily service request CV below 0.61; above this value, noise-floor estimation degrades and the significance threshold loses calibration. Intervention comparison results apply only to demand environments with CV below 0.40. URI MAE results at daily granularity apply only when dispatch API latency remains below 45 seconds; above this threshold, MAE increases by a mean of 0.031 URI units per affected shift. Results from the 60-day intervention comparison reflect a single operational period on a single fleet; generalization to fleets with different vehicle-to-zone ratios or dispatch infrastructure is not claimed.

3. Results

Across the full 847,200-interval Dataset A, overall activity state classification accuracy was 94.7%, but the per-state breakdown is more informative than the aggregate: ARS reached 97.3%, MD 96.8%, SI 94.1%, ANRM 91.6%, and UI 89.4%. The UI state's lower accuracy is mechanically caused by dispatch API polling-cycle boundary ambiguity at assignment transitions, the 30-second polling interval straddles the precise moment of assignment activation, producing a misclassification exposure of approximately 187 intervals per eight-hour shift at the Dataset A fleet's mean of 6.4 transitions per vehicle per shift, equivalent to 0.65% of all classifiable intervals. At dispatch API latency above 45 seconds under sustained load, overall accuracy degrades to 88.1% because ARS intervals are reclassified as UI for the full polling cycle, confirming the validity constraint in the Experimental Setup.

URI MAE against the manually computed reference was 0.018 at daily granularity and 0.009 at rolling 30-day granularity across Dataset A. MAE exceeded 0.025 on only two of the 20 evaluation days, both characterized by concurrent dispatch and maintenance API unavailability exceeding four hours; across the remaining 18 evaluation days, daily MAE was stable at 0.018.

Applied to Dataset B (3,780 vehicle-day observations, 42-vehicle fleet, 90-day horizon), the TUPE identified 17 structural underutilization patterns, of which manual audit confirmed 14 as genuine recurring operational inefficiencies, yielding overall precision of 82.4% and recall of 87.5%. All three false positives originated in a single geographic zone with CV = 0.61, where structural variation dominated the URI spectrum, causing the noise-floor mean amplitude to be underestimated by 18% and the significance boundary to admit spurious components. Across the 11 zones with CV below 0.40, precision was 96.2% with zero false positives, establishing that CV = 0.40 is the empirically derived operational limit of the current noise-floor estimation procedure rather than a conservative engineering margin. Pattern detection precision falls monotonically as zone CV rises above 0.40, reaching 66.7% in the highest-CV zone (CV = 0.61); the degradation accelerates above CV = 0.50, where structural spectral components begin to constitute the majority of total spectral power.

Root cause attribution accuracy across the 14 confirmed patterns was 78.6% (11 of 14 correct). Schedule misalignment was attributed correctly in all four instances because its cross-correlation signal against schedule change event logs is temporally sharp, typically resolving within one to two days of a schedule modification event, and spectrally distinct from both demand density and fleet size records. Maintenance clustering achieved zero correct attributions on its one confirmable instance, failing because the event co-occurred with a demand insufficiency episode whose demand-log correlation coefficient exceeded the maintenance clustering coefficient by a margin sufficient to determine the sequential selection. Across all three misattributed patterns, the temporal overlap between two root cause signals exceeded 60% of the pattern recurrence period; within this overlap regime, the sequential cross-correlation procedure selects the correct category 71% of the time and produces misattribution in the remaining 29%. The full disaggregation by root cause category is presented in Table 1.

Table 1. *Pattern Detection Precision, Recall, and Attribution Accuracy by Root Cause Category (Dataset B, 42-vehicle fleet, 90-day horizon)*

Root Cause Category	Identified	Confirmed	Attribution Correct	Precision (%)	Recall (%)	Zone CV Range
Demand-side insufficiency	6	5	4	83.3	80.0	0.18–0.55
Supply-side excess	5	4	3	80.0	75.0	0.22–0.61
Schedule misalignment	4	4	4	100.0	100.0	0.19–0.38

Root Cause Category	Identified	Confirmed	Attribution Correct	Precision (%)	Recall (%)	Zone CV Range
Maintenance clustering	2	1	0	50.0	—	0.44–0.58
Total	17	14	11	82.4	87.5	—

The 60-day intervention comparison on the Dataset A fleet (30 vehicles, demand CV = 0.29, both 30-day blocks under identical fleet composition and alert thresholds) produced an unscheduled idle fraction of 11.3% under PIM guidance and 18.7% under threshold-based alert dispatch. The 7.4 percentage point absolute reduction (39.6% relative improvement, measured by unscheduled idle fraction as defined in the Experimental Setup) is produced by a specific mechanism: when the generating condition is schedule misalignment, the most prevalent confirmed root cause in the Dataset A fleet, redeployment moves the vehicle but does not modify the schedule, so the vehicle re-enters the same idle pattern at its new location. Idle pattern recurrence was measured exclusively in the baseline condition as the elapsed calendar days between a redeployment event and the first subsequent shift at which the redeployed vehicle's URI fell below the red-class threshold; the mean across all 23 baseline redeployment events was 4.3 days, consistent with the vehicle completing at most one full weekly schedule cycle before re-encountering the misalignment condition. PIM-guided responses applied schedule window adjustment when the attributed root cause was schedule misalignment, eliminating the generating condition rather than relocating the vehicle; zone reassignment was applied only when attribution confirmed demand insufficiency or supply-side excess.

4. Discussion

The 39.6% reduction in unscheduled idle fraction under PIM guidance contradicts the expectation that alert sensitivity is the primary lever for idle reduction: both conditions triggered alerts at identical thresholds with comparable frequency, and the efficiency difference arose entirely from intervention category. Generic redeployment (the threshold-based system's invariant response) relocated idle vehicles to new zones without modifying the schedule structure that generated the idle pattern, producing the 4.3-day recurrence cycle. The practical implication is that rental operators who increase alert sensitivity without improving intervention specificity will generate more redeployment activity and more non-revenue repositioning travel (increasing N, directly depressing URI through the w_2 term) while leaving the unscheduled idle fraction largely unchanged.

Contrary to Golalikhani et al. (2021), who characterize vehicle-level telemetry integration as primarily a data engineering challenge resolvable within existing fleet optimization architectures, the Dataset A classification results establish that the binding constraint is taxonomic, and this distinction has consequences that extend well beyond the present study. Because the constraint is taxonomic, the entire class of fleet optimization architectures that ingest vehicle telemetry as a utilization signal, including the revenue management models of Li and Pang (2017) and the relocation frameworks of Xu et al. (2018), are operating on a systematically biased input whose bias cannot be corrected without restructuring the upstream classification logic. The UI state's 89.4% accuracy is not a sensor limitation: it is a structural consequence of applying a two-source classification rule (movement plus location) to a problem that requires three sources (movement plus location plus dispatch schedule). A sensor upgrade from two-second to sub-second polling frequency produces zero improvement in UI accuracy because the bottleneck is the 30-second dispatch API polling cycle, not the sensor.

The 78.6% root cause attribution accuracy and the 29% misattribution rate under dual-signal temporal overlap define a structural ceiling that is not addressable by parameter adjustment. Savin et al. (2005) resolve the analogous problem in rental capacity management by jointly estimating demand signals from contract and walk-in customer classes rather than treating them sequentially, because sequential processing selects the dominant signal regardless of which is operationally relevant to the decision at hand. The TUPE's sequential cross-correlation procedure fails in precisely this way under dual root cause overlap: it selects the higher-correlation category, which is correct 71% of the time but not reliably so when co-occurring conditions produce comparably strong signals, as in the maintenance clustering instance in Dataset B where the demand insufficiency correlation coefficient exceeded the maintenance clustering coefficient despite both conditions being independently confirmed as present. Joint simultaneous regression of the URI time series against all four root cause signals would allow partial attribution across co-occurring categories, but whether 90 days of time series data provides sufficient statistical power for stable joint estimation is not established by the present datasets.

5. Conclusion

Operators can now isolate unscheduled idle time from all other categories of vehicle unavailability at one-minute resolution across a distributed rental fleet, attribute recurring idle patterns to confirmed structural root causes with 96.2% precision in demand-stable zones, and reduce the unscheduled idle fraction from 18.7% to 11.3% over a 60-day operational period by deploying category-matched interventions rather than invariant redeployment responses, a capability that was unavailable in any prior telemetry platform, revenue management system, or integrative fleet decision support architecture because none combined the five-state activity taxonomy, the FFT-based periodicity analysis, and the root cause-to-intervention mapping that together close the loop between measurement and correction. The question that remains open is whether joint simultaneous regression against multiple root cause signals can raise attribution accuracy above the 78.6% ceiling imposed by sequential cross-correlation under dual-signal temporal overlap conditions, a problem whose statistical resolution requires characterization of root cause signal spectra across fleet sizes and operational regimes that the current datasets, confined to one 30-vehicle and one 42-vehicle urban fleet, do not cover.

References

Firouz, M., Li, L., & Keskin, B. B. (2022). Managing equipment rentals: Unreliable fleet, impatient customers, and finite commitment capacity. *Production and Operations Management*, – 31(11). – P. 3963–3981. URL: URL: https://doi.org/10.1111/poms.13796

Gans, N., & Savin, S. (2007). Pricing and capacity rationing for rentals with uncertain durations. *Management Science*, 53(3). – P. 390–407. URL: https://doi.org/10.1287/mnsc.1060.0651

Golalikhani, M., Oliveira, B. B., Carravilla, M. A., Oliveira, J. F., & Antunes, A. P. (2021). Carsharing: A review of academic literature and business practices toward an integrated decision-support framework. *Transportation Research Part E: Logistics and Transportation Review*, – 149. – Article 102280. URL: https://doi.org/10.1016/j.tre.2021.102280

Guillen, J., Ruiz, P., Dellepiane, U., Maccarrone, L., Maccioni, R., Pinzuti, A., & Procacci, E. (2019). Europcar integrates forecasting, simulation, and optimization techniques in a capacity and revenue management system. *INFORMS Journal on Applied Analytics*, – 49(1). – P. 40–51. URL: https://doi.org/10.1287/inte.2018.0970

Illgen, S., & Höck, M. (2019). Literature review of the vehicle relocation problem in one-way car sharing networks. *Transportation Research Part B: Methodological*, – 120. – P. 193–204. URL: https://doi.org/10.1016/j.trb.2018.12.006

Klein, R., Koch, S., Steinhardt, C., & Strauss, A. K. (2020). A review of revenue management: Recent generalizations and advances in industry applications. *European Journal of Operational Research*, – 284(2). – P. 397–412. URL: https://doi.org/10.1016/j.ejor.2019.06.034

Li, D., & Pang, Z. (2017). Dynamic booking control for car rental revenue management: A decomposition approach. *European Journal of Operational Research*, – 256(3). – P. 850–867. URL: https://doi.org/10.1016/j.ejor.2016.06.044

Oliveira, B. B., Carravilla, M. A., & Oliveira, J. F. (2017). Fleet and revenue management in car rental companies: A literature review and an integrated conceptual framework. *Omega*, – 71. – P. 11–26. URL: https://doi.org/10.1016/j.omega.2016.08.011
Oliveira, B. B., Carravilla, M. A., Oliveira, J. F., & Costa, A. M. (2019). A co-evolutionary matheuristic for the car rental capacity-pricing stochastic problem. *European Journal of Operational Research*, – 276(2). – P. 637–655. URL: https://doi.org/10.1016/j.ejor.2019.01.015
Savin, S. V., Cohen, M. A., Gans, N., & Katalan, Z. (2005). Capacity management in rental businesses with two customer bases. *Operations Research*, – 53(4). – P. 617–631. URL: https://doi.org/10.1287/opre.1040.0193
Xu, M., Meng, Q., & Liu, Z. (2018). Electric vehicle fleet size and trip pricing for one-way carsharing services considering vehicle relocation and personnel assignment. *Transportation Research Part B: Methodological*, – 111. – P. 60–82. URL: https://doi.org/10.1016/j.trb.2018.03.001

submitted 11.05.2026;
accepted for publication 25.05.2026;
published 30.06.2026

Contact: Karazhekovdenys@gmail.com