Poster Abstract: Robustness of Smart Parking Assignment Under Realistic Sensor Noise

Behafarid Hemmatpour
IMDEA Networks Institute
Universidad Carlos III de Madrid
Madrid, Spain
behafarid.hemmatpour@networks.imdea.org

Javad Dogani
IMDEA Networks Institute
Madrid, Spain
javad.dogani@networks.imdea.org

Nikolaos Laoutaris
IMDEA Networks Institute
Madrid, Spain
nikolaos.laoutaris@networks.imdea.org

***Abstract*—Searching for on-street parking remains a persistent urban challenge. Building on our prior coordinated parking framework (*Cord-Approx*), we upgrade its predictive engine with an attention-based multi-horizon model to better capture dynamic availability. We then stress-test the system under realistic imperfect sensing by introducing reduced sensor coverage (ρ) and false vacancy rate (ϕ), and show that, in the evaluated setting, coordinated assignment retains a clear advantage under noisy sensing. Even under severe noise, with 40% of free-spot detections missing (ρ=0.6) and a 5% false vacancy rate (ϕ=0.05), our system delivers a parking success ratio 2.3× that of competitors (non-users) and 51.9% lower search time.**

I. INTRODUCTION AND MOTIVATION

In dense metropolitan centers, searching for on-street parking is a major contributor to traffic congestion. In our foundational work [1], we established the benefits of coordinated on-street parking allocation, showing that leveraging historical data to estimate competitor behavior (*Cord-Approx*) outperforms uncoordinated greedy search (*Unc-Agn*) under ideal sensor conditions. To move toward deployable municipal systems, coordination must remain effective under imperfect sensing. Real-world urban sensing systems, such as crowdsourced deployments [2], often face incomplete spatial coverage and erroneous vacancy reports. When sensing is noisy, guidance based only on reported vacancies can backfire by routing drivers to occupied spaces and increasing search time. In this work, we show that, in the evaluated setting, coordination plus prediction retains an advantage as sensor coverage ρ decreases and false vacancy rate ϕ increases. We integrate a multi-horizon predictor into *Cord-Approx* and quantify a practical robustness boundary over the evaluated (ρ, ϕ) settings.

II. NOISE MODEL & PREDICTIVE ARCHITECTURE

A common simplification in smart-parking simulations is treating sensor uncertainty as bit-flip noise (e.g., reporting an occupied spot as "free"). This can inject *phantom capacity*, i.e., the perceived free set can exceed physical capacity, presenting the optimization algorithm with more spatial resources than physically exist. To avoid this, we use a **replacement-based sensor noise model**. The system's perceived state is constrained by two physical parameters:

- **Sensor Coverage** ($\rho \in [0, 1]$)**:** The probability that a physically free spot is detected as free.

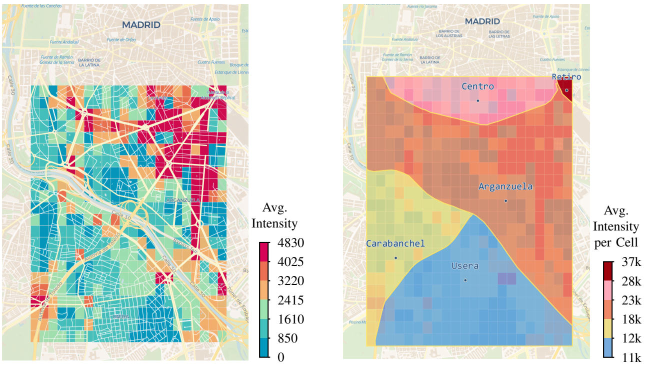

(a) (b)

Fig. 1. Spatial distribution of traffic intensity in the central Madrid study area during peak hours (09:00–17:00). (a) Fine-grained traffic intensity map at the grid-cell level, showing strong spatial heterogeneity across the study area [1]. (b) District-level aggregation of the same intensity signal, where each value represents the average grid-cell intensity within a district. The underlying grid uses 7-character Geohash cells ($\approx 152.8\,\mathrm{m} \times 116.4\,\mathrm{m}$).

- **False Vacancy Rate** ($\phi \in [0, 1]$)**:** The false positive rate where an occupied spot erroneously signals vacancy.

False vacancies actively *replace* valid detections, ensuring the system never perceives more capacity than reality allows. To operate under these sensing imperfections, we upgrade *Cord-Approx* by integrating an attention-based multi-horizon predictor (PatchTST-lite). The predictor produces an arrival-aware reliability score for each candidate spot/cell based on predicted short-term dynamics. We inject this signal into the Hungarian algorithm cost matrix as an additional penalty, so spots reported as "free" by sensors but predicted to be volatile receive a higher cost and are less likely to be selected. This reduces participants' exposure to assignments driven by false vacancies.

III. EXPERIMENTAL SETUP

We evaluate our framework using a high-fidelity, discrete-event simulation of the central Madrid study area (Fig. 1, generated with kepler.gl[1]). Grounded in over 2 million municipal sensor reports [3], the environment manages 12,365 physical parking spots across a 22×22 grid of 7-character Geohash

---

[1] Basemap uses CARTO Voyager; map data © OpenStreetMap contributors.

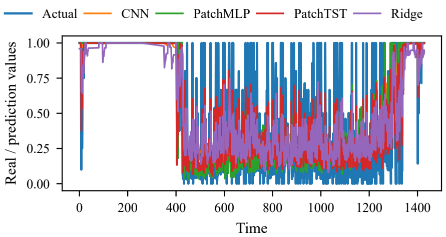

Fig. 2. Multi-horizon prediction example: PatchTST-lite avoids spurious spikes compared to CNN, PatchMLP, and Ridge baselines (cell 240 shown).

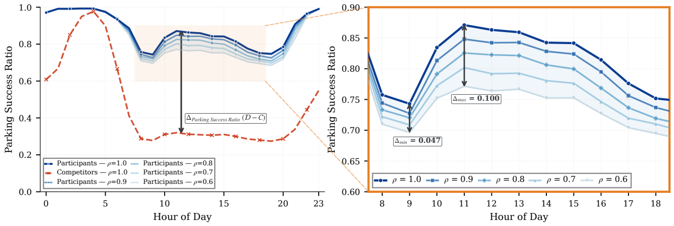

Fig. 3. Robustness of participants' parking success ratio under degraded ρ, with a zoomed view of peak hours (09:00–17:00).

cells [4]. We simulate a mixed-traffic ecosystem with two agent classes: **participants** (app-guided users) and **competitors** (uninformed non-users using greedy visual search).

IV. EVALUATION HIGHLIGHTS

We evaluate our system's operational boundaries via a 24-hour Madrid simulation. Under perfect sensing (ρ=1.0, ϕ=0.0), our prior *Cord-Approx* system achieved 77.54% parking success ratio [1]. With PatchTST-lite, we raise this to 82.2%, narrowing the gap to the *Cord-Oracle* upper bound (85.32%) [1]. Under sensor degradation, we observe a clear asymmetry: performance degrades more gradually with reduced ρ than with increasing ϕ. Even under severe noise (ρ=0.6, ϕ=0.05), participants still find parking spots in less than half the time of competitors (Table I), while maintaining more than double their parking success ratio (Fig. 4).

Predictive Accuracy Upgrade: Building on our prior *Cord-Approx* framework [1], we upgrade its prediction engine to better track short-term availability. Fig. 2 illustrates the improved stability: PatchTST-lite avoids spurious spikes during rapid fluctuations compared to alternative baselines.

Robustness to Sensor Coverage (ρ): Performance degrades gradually as sensor coverage decreases. As shown in Fig. 3, when coverage drops from ρ=1.0 to ρ=0.6, the peak-hour separation between the participants success ratio curves remains bounded: the worst-case drop is $\Delta_{\max}$=0.10 in success ratio, i.e., 10 percentage points (pp). Even with 40% reduced sensor coverage, participants still achieve more than twice the parking success ratio of competitors during peak hours.

Penalty of Misinformation: A higher false vacancy rate (ϕ) degrades the parking success ratio more sharply than reduced sensor coverage (ρ). False vacancies trap participants in rerouting after phantom assignments, inflicting severe time penalties. In Table I at a mid-range operating point (ρ=0.8, ϕ=0.08), participants average 8.49 min versus 19.76 min for competitors, i.e., 11.27 min less ($\sim$57% reduction). However, Table I shows that even under degraded settings (ρ=0.6, ϕ=0.05), the average participant search time is only 9.42 min, still cutting the nearly 20-minute search time (19.57 min) of competitors by more than half. Fig. 4 shows the evaluated ρ, ϕ settings. Even under sparse sensor coverage ($\rho = 0.6$) and high noise ($\phi = 0.15$), participants sustain more than double the parking success ratio of competitors, suggesting $\phi \approx 0.15$ as a practical robustness boundary in the evaluated settings.

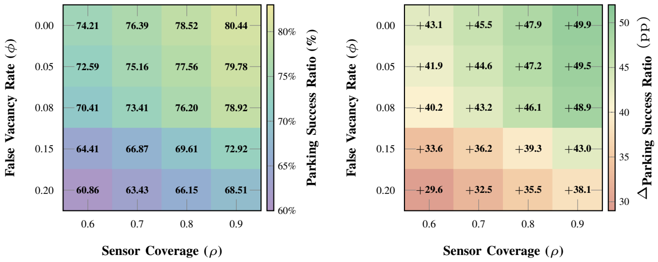

(a) (b)

Fig. 4. Peak-hour robustness across $\rho \times \phi$ under imperfect sensing. (a) Average participants' parking success ratio as sensor coverage ρ decreases and false vacancy rate ϕ increases. (b) participants–competitors parking success ratio gap in percentage points (pp) across evaluated noise settings.

TABLE I
AVERAGE SEARCH TIME (MIN) FOR SELECTED SETTINGS.

Agent	ρ=0.6 ϕ=0.05	ρ=0.7 ϕ=0.08	ρ=0.8 ϕ=0.08	ρ=0.9 ϕ=0.15	ρ=1.0 ϕ=0.00
Participants	9.42	9.22	8.49	9.51	6.85
Competitors	19.57	19.72	19.76	19.76	19.70

V. TAKEAWAY

To transition from theory to civic deployment, municipal IoT-enabled smart-parking systems must account for imperfect sensing. Our stress-test shows an asymmetric robustness trade-off: performance is relatively robust to reduced sensor coverage, but degrades more sharply as false vacancy rate increases (Fig. 4). Accordingly, in our evaluation, improving sensing accuracy (lower ϕ) matters more than marginally expanding coverage (higher ρ), while coordination plus prediction preserves advantages over competitors across the evaluated (ρ, ϕ).

REFERENCES

[1] B. Hemmatpour, J. Dogani, and N. Laoutaris, "Reducing street parking search time via smart assignment strategies," in *Proceedings of the 33rd ACM International Conference on Advances in Geographic Information Systems*, pp. 378–381, 2025.

[2] D. Zhao, Z. Cao, C. Ju, D. Zhang, and H. Ma, "D2park: Diversified demand-aware on-street parking guidance," in *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies (IMWUT, Vol. 4)*, vol. 4 of *IMWUT '20*, pp. 1–25, Association for Computing Machinery, 2020.

[3] Madrid City Council, "Historical traffic data." https://bit.ly/43xk6sn, 2024. Accessed: 2024-04-18.

[4] Geohash.es, "Geohash.es — geohash encoding and decoding."