Quantifying the Robustness of Smart Street Parking Assignment to Sensor Noise

1st Behafarid Hemmatpour
IMDEA Networks Institute
Universidad Carlos III de Madrid
Madrid, Spain
behafarid.hemmatpour@networks.imdea.org

2nd Javad Dogani
IMDEA Networks Institute
Madrid, Spain
javad.dogani@networks.imdea.org

3rd Nikolaos Laoutaris
IMDEA Networks Institute
Madrid, Spain
nikolaos.laoutaris@networks.imdea.org

***Abstract*—Searching for street parking remains a major source of congestion in dense cities. Prior studies either assume perfect sensing of free spots or study sensing separately, leaving the robustness of smart parking assignment under sensing errors unclear. We address this using a Madrid-calibrated city-scale simulator and a replacement-based sensor-noise model parameterized by sensor coverage ρ and false vacancy rate ϕ, where missed detections hide real free spots and false vacancies report occupied spots as free. We extend our *Cord-Approx* strategy by embedding arrival-aware predictions into the Hungarian assignment cost and by rerouting participants after false vacancy encounters. Under perfect sensing, the parking success ratio rises to 82.2%, up from 77.54% in our prior work and closer to the 85.32% *Cord-Oracle* upper bound. Even at ρ=0.6, ϕ=0, participants still achieve 74.2% success and 8.92 minutes of search time, versus 31.1% and 19.45 minutes for competitors. This coordinated advantage remains visible in all five Madrid districts in a representative noisy setting and remains positive across the evaluated grid, with a practical operating boundary around $\phi \approx 0.15$.**

I. Introduction

The management of on-street parking remains a critical challenge in urban planning. In dense metropolitan centers, limited on-street parking capacity drives fierce competition, leading to "cruising for parking", a phenomenon estimated to account for up to 30% of traffic in some high-demand districts [1]. While intelligent parking systems leveraging Internet of Things (IoT) sensors and cloud analytics promise to reduce search time by guiding drivers to available spots, a critical gap remains between theoretical models and real-world deployment. Most existing studies, including our own prior work [2], operate under the premise of *perfect sensing*, assuming the central coordinator possesses real-time, error-free knowledge of every parking spot. In practice, deployments operate with incomplete coverage and sensing errors rather than perfect observability [3], so robustness must be evaluated under partial and unreliable system views.

However, robustness analysis is meaningful if sensing errors are modeled in a physically valid way. A critical challenge in smart parking research is how to model sensing errors without violating the physical conservation of parking capacity. Several smart-parking studies represent imperfect sensing through binary misclassification with false positives (false vacancies) and false negatives (missed vacancies), i.e., probabilistic errors on occupancy observations [4], [5]. If false vacancies are implemented additively, the perceived free-spot set can exceed the physical supply, creating phantom capacity and artificially relaxing the assignment problem. To avoid this inflation paradox, we introduce a **Replacement-Based Sensor Noise Model** in which false vacancies *replace* a subset of valid detections rather than being appended to them. This preserves the conservation of physical space while forcing the dispatcher to operate under genuinely degraded information. Intuitively, the model swaps some correctly observed free spots with false-vacancy reports, rather than adding phantom spots on top of the real supply.

We extend our *Cord-Approx* strategy with a robust dispatcher that uses an arrival-aware success estimate and explicit phantom-spot handling. Unlike our prior four-strategy comparison, this extension focuses on the deployable *Cord-Approx* strategy and stress-tests it under imperfect sensing. *Cord-Approx* approximates competitor effects via learned spatiotemporal success probabilities rather than tracking competitor locations. While our previous work [2] established the baseline simulator and predictive engine, this extension fundamentally re-engineers the assignment logic to operate under the **Replacement-Based Sensor Noise Model**.

Our innovation is to turn the predictive layer into a reliability filter for sensor-reported free spots under (ρ, ϕ). Using a lightweight attention-based multi-horizon model, the system computes an arrival-aware success estimate for each candidate spot and embeds it into the Hungarian cost matrix to penalize assignments that are unlikely to remain feasible upon arrival. The dispatcher also accounts for phantom-spot encounters by re-queuing failed assignments at the next decision step. Thus, the core contribution is a capacity-preserving robustness framework that couples noisy sensing, arrival-aware prediction, and rerouting in the deployable *Cord-Approx* pipeline; the robustness analysis under imperfect sensing is developed in the methodology and results sections below.

A. Contributions

This work provides a city-scale sensitivity analysis of on-street parking coordination under realistic, non-inflationary sensor noise. Using city-scale traffic intensity data from Madrid, Spain [6], we simulate the interaction between par-

ticipants (app users) and competitors (non-users). Our specific contributions are:

- We formalize a replacement-based sensor noise model, defined by sensor coverage ρ and false vacancy rate ϕ, that prevents capacity inflation and enables realistic stress testing of parking assignment under incomplete and misleading sensing.
- We propose an attention-based multi-horizon forecasting model that produces arrival-aware success probabilities for future parking outcomes under temporal variation and sensing uncertainty.
- We extend *Cord-Approx* with a robust dispatcher that integrates these arrival-aware success probabilities into the Hungarian assignment cost and explicitly handles phantom-spot rerouting, enabling city-scale robustness analysis under imperfect sensing.

The remainder of this paper is organized as follows: Section II reviews the literature. Section III defines the system model. Section IV details the robust assignment algorithm and noise models. Section V describes the experimental setup using the Madrid dataset, and Section VI presents the results.

II. Related Work

We organize prior work into four themes: (i) urban impact of parking search, (ii) sensing and availability prediction, (iii) coordinated allocation, and (iv) robustness under imperfect information.

A. Urban Impact and Searching Dynamics

The foundational understanding of parking search behavior stems from Shoup [1], who characterized curbside searching as a communal inefficiency that accounts for approximately 30% of traffic in some high-demand districts. Fosgerau and de Palma [7] further formalized how parking conditions interact with urban traffic congestion, reinforcing that parking search imposes system-level travel and delay externalities. Recent studies have specialized these findings; for example, Dalla Chiara et al. [8] demonstrated that curb availability information specifically reduces search time for delivery drivers, highlighting the logistical benefits of information-centric systems.

B. Sensing Infrastructure and Availability Prediction

Visibility is necessary for guidance. Early sensing research utilized probe vehicles in the ParkNet system [4] and smartphone-based Wi-Fi sensing in ParkSense [9] to map occupancy. Recent efforts leverage multiple data streams: Inam et al. [10] integrated occupancy, weather, and traffic data using Random Forest models to predict block-level availability. Similarly, Zhao et al. [11] proposed D2Park, which minimizes cruising costs via diversified demand-aware guidance. While sensing via satellite imagery [12] and fixed/mobile hybrid systems [13] have improved coverage, these works primarily facilitate *guidance* for independent agents rather than system-wide coordination, and they typically do not address robustness when sensing is incomplete or unreliable.

C. Coordinated Allocation: From Theory to Real-World Data

To resolve competition, the literature shifted toward centralized assignment. Geng and Cassandras [14] theoretically proved that coordinated allocation can outperform greedy visual search under their model assumptions. Game-theoretic approaches, such as those by Tan et al. [15], validated that reservation-based models provide better stability than competitive searching under fluctuating demand. Other systems, like iParker [16], utilized the Hungarian Algorithm [17] to manage dynamic resource allocation. Many coordination studies evaluate under simplified demand assumptions or limited-scale settings; our prior work [2] moved to a city-scale, data-driven calibration using Madrid traffic intensity. That study demonstrated that when grounded in real-world geospatial constraints and traffic flows, the Coordinated Approximate strategy (*Cord-Approx*) can reduce search times by 66% compared to the uncoordinated baseline.

D. Robustness to Information Uncertainty

Despite the empirical success of data-driven coordination [2], a critical dependency remains: the assumption of perfect sensing (i.e., $\rho = 1, \phi = 0$). Evaluation of the SFpark pricing experiment [18], [3] revealed that physical sensor reliability remains a bottleneck in deployment. Rodríguez et al. [19] analyzed user behavior under different parking-information scenarios provided by smart devices and connected cars, while Higuchi et al. [20] analyzed how parking preferences are affected by inconsistent data. A common modeling shortcut is to treat sensing uncertainty as independent per-spot misclassification with fixed false-positive/false-negative rates (i.e., "flips" in reported occupancy) [4], [5]. If false vacancies are implemented additively (without a compensating removal of true detections), the perceived free-spot set can exceed the physical supply ("phantom capacity"), which may artificially improve simulated performance by relaxing capacity constraints.

E. Positioning and Contribution

We build directly on the data-driven foundation established in [2]. Whereas our previous work studied coordinated assignment under ideal sensing, this paper focuses on robustness under sensing degradation. Specifically, we replace the idealized sensor view with a replacement-based noise model, upgrade the predictive layer to a multi-horizon arrival-aware success estimator, and quantify when coordinated assignment remains operational under (ρ, ϕ) using the Madrid dataset [6]. Our contribution is threefold. First, we replace the idealized sensor assumptions of [2] with a **replacement-based sensor noise model**, which prevents the artificial capacity inflation common in additive noise models. Second, we transform the predictive engine from an availability forecaster into a *reliability filter* using an attention-based multi-horizon model (PatchTST-lite). Third, we provide a city-scale quantification of the rerouting penalty induced by phantom-spot encounters and identify a practical operating boundary, defined by the participants–competitors performance gap, beyond which coordinated gains weaken substantially. Since external systems use different data,

TABLE I
CORE NOTATION USED IN THE SYSTEM MODEL, METHODOLOGY, AND EXPERIMENTAL SETUP.

Symbol	Definition
Simulation & Grid	
$N \times N$	Grid dimension (22×22 Geohash cells).
T, t	Simulation horizon and time step ($t \in \{1, \dots, T\}$).
$t_{\max}$	Maximum search time before timeout (30 min).
$\mathcal{S}$	Set of all parking spots in the city.
$\mathcal{S}_t$	Physically available (ground-truth) spots at t.
Agents	
$\mathcal{D}$, D_t	All participants; subset active at t.
$\mathcal{C}$, C_t	All competitors; subset active at t.
$\tau(d,s)$, $\tau(c,s)$	Manhattan distance (travel time) from participant d or competitor c to spot s.
R	Competitor observability radius (in grid cells).
$\mathcal{V}_{c,t}$	Set of physically available spots visible to competitor c at time t.
Sensor Noise Model	
ρ	Detection probability of a truly free spot (sensor coverage).
ϕ	False vacancy rate (probability of a phantom-spot reading).
$\mathcal{S}_{valid}$	Subset of truly free spots correctly detected after coverage filtering.
$\mathcal{G}$	Generated phantom set (from occupied spots).
$\hat{\mathcal{S}}_t$	Perceived spots at t ($\hat{\mathcal{S}}_t \neq \mathcal{S}_t$).
Cost & Prediction	
$\mathbf{M}_t$	Assignment cost matrix at t.
$\mathbf{M}_t(d,s)$	Entry (d,s) of the cost matrix $\mathbf{M}_t$.
$x_{d,s,t}$	Binary assignment variable; equals 1 if participant d is assigned to perceived spot s at time t.
$\mathcal{A}_t^{partc}$	All spots assigned to participants at time t.
$p_{k,t}$	Observed parking success ratio in cell k at forecasting tick t.
$\hat{\mathbf{p}}_{k,t}$	Vector of multi-horizon parking success-ratio predictions.
$\hat{p}_{k,t}^{(h)}$	Predicted parking success ratio for cell k at horizon h.
$k(s)$	Mapping from spot s to its Geohash index k.
$h(d,s)$	Arrival-aligned horizon index for pair (d,s).
$p^{\text{arr}}(d,s)$	Arrival-aware parking success estimate.
$\tilde{\tau}(d,s)$	Effective distance used as the robust assignment cost.
H	Number of forecast horizons (ticks)
L	Input context window length (ticks)
Δt_f	Forecasting tick size (e.g., 1 min)

sensing assumptions, and simulator settings, we use our prior Madrid-calibrated *Cord-Approx* study as the direct numerical baseline and use other coordinated and learning-based systems for positioning rather than claiming a cross-system ranking.

III. SYSTEM MODEL AND PROBLEM FORMULATION

We consider a dynamic urban parking scenario on a grid over a discrete horizon $t \in T$, with one-minute time steps. The ecosystem consists of two distinct agent classes competing for a finite set of on-street parking resources.

Let $\mathcal{S}$ be the set of all on-street parking spots in the city. At any time step t, the set of physically available spots is $\mathcal{S}_t \subseteq \mathcal{S}$, while the coordinator acts on a perceived set $\hat{\mathcal{S}}_t$.

• **Participants** ($D_t \subseteq \mathcal{D}$): Participants are drivers equipped with our smart parking application. Upon entering the system, they request a parking assignment and follow the system's guidance. They do not search visually; they rely entirely on the system to provide a target spot s. If a participant reaches an assigned spot that is physically occupied (a phantom-spot encounter), the system records the encounter and the participant becomes eligible for reassignment at the next step (Algorithm 1). The dispatcher recomputes assignments each step; agents not successfully parked remain active at $t+1$.

• **Competitors** ($C_t \subseteq \mathcal{C}$): Competitors represent the background traffic (non-users) who search for parking visually. They are uncoordinated and do not share their destination or position with the system. A competitor $c \in C_t$ detects and claims any spot $s \in \mathcal{S}_t$ within a visual observability radius R measured in Manhattan distance. The set of spots visible to competitor c at time t is:

$$\mathcal{V}_{c,t} = \{s \in \mathcal{S}_t \mid \tau(c,s) \leq R\} \tag{1}$$

where $\tau(c,s)$ is the travel time between competitor c and spot s, and competitors are unaffected by system sensing because they detect spots directly within radius R.

A critical challenge addressed in this paper is that the perceived set $\hat{\mathcal{S}}_t$ differs from the ground truth $\mathcal{S}_t$. The system may be blind to real spots (False Negatives, controlled by ρ) or hallucinate free spots that are actually occupied (False Positives; controlled by ϕ). In this work, we specifically focus on the *Cord-Approx* strategy. Unlike full-information oracle settings, *Cord-Approx* optimizes assignments using only (1) the current positions of participants, (2) a *perceived* set of parking spots $\hat{\mathcal{S}}_t$ derived from imperfect sensors, and (3) historical context to estimate arrival-aware parking success for robust assignment. In this extension, the coordinator optimizes over the perceived set $\hat{\mathcal{S}}_t$, while the simulator resolves the physical outcome against the ground-truth set $\mathcal{S}_t$. The objective is to assign each active participant $d \in D_t$ to at most one perceived spot $s \in \hat{\mathcal{S}}_t$ so as to minimize the total assignment cost:

$$\min_{x_t} \sum_{d \in D_t} \sum_{s \in \hat{\mathcal{S}}_t} \mathbf{M}_t(d,s)\, x_{d,s,t} \tag{2}$$

subject to

$$\sum_{s \in \hat{\mathcal{S}}_t} x_{d,s,t} \leq 1, \qquad \forall d \in D_t, \tag{3}$$

$$\sum_{d \in D_t} x_{d,s,t} \leq 1, \qquad \forall s \in \hat{\mathcal{S}}_t, \tag{4}$$

where $x_{d,s,t} \in \{0,1\}$ and $\mathbf{M}_t(d,s)$ is the assignment cost defined in Section IV. This matching is recomputed each step using current positions and the perceived set $\hat{\mathcal{S}}_t$; physical outcomes, including competitor preemption and phantom-spot encounters, are resolved by the simulator dynamics described in Section IV. Consequently, each dispatch step is solved exactly for the current perceived one-step cost matrix, but end-to-end performance guarantees are empirical because future arrivals, competitor preemption, and sensing errors are stochastic. Table I summarizes the key symbols used in the model and methodology throughout the paper.

IV. METHODOLOGY

We extend *Cord-Approx* with a robust coordinated assignment framework for imperfect sensing. The methodology has four components: (i) constructing a replacement-based perceived set $\hat{\mathcal{S}}_t$ from the ground-truth free-spot set $\mathcal{S}_t$, (ii) training an attention-based multi-horizon forecasting model that predicts cell-level future parking success ratios, (iii) converting these forecasts into an arrival-aware success estimate, and (iv) executing runtime dispatch with explicit phantom-spot handling. The forecasting model is trained offline, whereas the runtime dispatcher operates in three phases: sensor-noise generation, cost-matrix construction, and dispatch/rerouting.

A. Replacement-Based Sensor Noise Model

Real-world sensor networks suffer from two distinct error types: *False Negatives* (FN), where available spots are missed due to coverage gaps or packet loss; and *False Positives* (FP), where occupied spots are erroneously reported as free (e.g., sensor calibration drift). To model this realistically without artificially inflating the system's perceived capacity, we introduce a **replacement protocol**. The system's view $\hat{\mathcal{S}}_t$ is constructed from the ground truth $\mathcal{S}_t$ using two parameters:

- $\rho \in [0, 1]$: **Sensor coverage**. This represents the probability that a physically available spot is successfully observed by the system. A value of $\rho < 1.0$ simulates coverage gaps where the system is "blind" to existing inventory (Omission).
- $\phi \in [0, 1]$: **False vacancy rate**. This represents the probability that a physically occupied spot erroneously signals vacancy due to sensor noise. These false vacancies constitute active misinformation (a commission error).

The perceived set $\hat{\mathcal{S}}_t$ is generated in three phases at each t:

1) We first identify the subset of correctly detected real spots, $\mathcal{S}_{valid} \subseteq \mathcal{S}_t$. Each spot in $\mathcal{S}_t$ is included in $\mathcal{S}_{valid}$ with probability ρ. This simulates the system being "blind" to a portion of the inventory.
2) We identify a set of potential false vacancies $\mathcal{G} \subseteq (\mathcal{S} \backslash \mathcal{S}_t)$. Each occupied spot $s \in (\mathcal{S} \backslash \mathcal{S}_t)$ generates a false-vacancy signal with probability ϕ.
3) To preserve physical capacity, we do not append $\mathcal{G}$ to $\mathcal{S}_{valid}$. Instead, let

$$m = \min\big(|\mathcal{G}|,\ |\mathcal{S}_{valid}|\big).$$

We then replace m uniformly sampled spots from $\mathcal{S}_{valid}$ with m uniformly sampled phantom spots from $\mathcal{G}$.

By construction, $|\hat{\mathcal{S}}_t| = |\mathcal{S}_{valid}|$ (replacement keeps the perceived-free count unchanged) and $|\mathcal{S}_{valid}| \leq |\mathcal{S}_t|$, so the system never perceives more free spots than physically exist. However, as ϕ increases, the reliability of the reported spots degrades because a larger fraction of perceived vacancies are phantom spots. We therefore use the replacement model as a conservative, capacity-preserving abstraction for stress-testing sensing errors, not as the only possible error process; spatially or temporally correlated failures are left to future work.

B. Attention-Based Multi-Horizon Forecasting

This subsection details the estimation of cell-level future parking success-ratio forecasts with a lightweight *attention-based multi-horizon regressor*. The model (i) captures diurnal and weekly patterns together with temporal regime shifts, and (ii) outputs *arrival-aware* success-ratio forecasts that integrate directly into the assignment cost defined in Phase 2 of Algorithm 1. Unlike an availability classifier, the objective here is to forecast, for each cell, the expected parking success ratio at future time steps.

1) Prediction Task and Target Definition: Let $k \in \{1, \ldots, K\}$ index spatial cells (Geohash grid), and let time be discretized into forecasting ticks of size Δt_f (e.g., $\Delta t_f = 1$ minute). For each cell k and tick t, let $p_{k,t} \in [0, 1]$ denote the observed parking success ratio at that cell and time step. If the simulator records counts, this can be defined as

$$p_{k,t} \triangleq \frac{n^{\mathrm{succ}}_{k,t}}{\max\Big(1, n^{\mathrm{att}}_{k,t}\Big)}, \tag{5}$$

where $n^{\mathrm{att}}_{k,t}$ is the number of parking attempts (or candidate arrivals) associated with cell k at tick t, and $n^{\mathrm{succ}}_{k,t}$ is the number of successful parking outcomes among them.

At each tick t, the model predicts a *multi-horizon* success-ratio vector

$$\hat{\mathbf{p}}_{k,t} = \big(\hat{p}^{(1)}_{k,t}, \ldots, \hat{p}^{(H)}_{k,t}\big), \qquad \hat{p}^{(h)}_{k,t} \approx p_{k,t+h}, \tag{6}$$

where H is the prediction horizon in ticks (e.g., $H = \lceil t_{\max}/\Delta t_f \rceil$, where $t_{\max}$ denotes the maximum search time). This formulation is therefore a *multi-horizon regression* task over a bounded target in $[0, 1]$, rather than a binary classification task.

2) Input Representation: For each cell k and tick t, the model consumes a fixed-length context window of length L ticks, denoted as

$$\mathbf{X}_{k,t}, \qquad \mathbf{X}_{k,t}[\ell, :] = \mathbf{x}_{k,t-L+\ell}, \quad \ell = 1, \ldots, L. \tag{7}$$

The feature vector $\mathbf{x}_{k,t}$ includes:

- **Historical success-ratio signal:** $p_{k,t} \in [0, 1]$, i.e., the observed parking success ratio of cell k at tick t;
- **Calendar covariates:** hour-of-day encoded by $\sin / \cos$, and day-of-week encoded by $\sin / \cos$ or one-hot features;
- **Static cell identity:** a learned embedding $\mathbf{e}_k$ that represents the spatial identity of cell k.

Continuous features are standardized by z-normalization. Missing values are imputed using the last observed value, and a binary indicator is added to mark imputed entries.

3) Model: PatchTST-lite with Attention: We adopt a patch-based Transformer encoder ("PatchTST-lite"). Given $\mathbf{X}_{k,t}$, we partition the input into $\mathcal{M}$ non-overlapping patches of length P. Each patch is projected into a latent space of dimension d_{model}, combined with positional encodings and the cell embedding, and then processed by N_{enc} Transformer encoder blocks. We apply mean pooling across tokens and use a regression head to produce an H-dimensional output vector

$\mathbf{u}_{k,t}$. To keep predictions in the valid range, we map it through an element-wise sigmoid:

$$\hat{\mathbf{p}}_{k,t} = \sigma(\mathbf{u}_{k,t}). \tag{8}$$

4) Training and Integration: The model is trained using a multi-horizon regression loss

$$\mathcal{L} = \frac{1}{|\mathcal{B}|H} \sum_{(k,t)\in\mathcal{B}} \sum_{h=1}^{H} \mathrm{Huber}\left(\hat{p}_{k,t}^{(h)},\, p_{k,t+h}\right),$$

where $\mathcal{B}$ is the mini-batch and $\mathrm{Huber}(\cdot,\cdot)$ denotes the standard Huber regression loss, which is quadratic for small residuals and linear for large residuals. We optimize the model with AdamW and cosine learning-rate decay.

At assignment time, for a candidate pair (d, s) with s in cell $k(s)$, the estimated travel time $\tau(d, s)$ is mapped to a horizon index

$$h(d, s) = \mathrm{clip}(\lceil \tau(d, s)/\Delta t_f \rceil\,,\ 1,\ H) \tag{9}$$

and the arrival-aware success ratio is defined as

$$p^{\mathrm{arr}}(d, s) = \hat{p}_{k(s),t}^{(h(d,s))}. \tag{10}$$

This value is used as the predictive reliability term in the assignment cost, so spots in cells with lower expected success ratio at the estimated arrival time receive a larger penalty.

C. Predictive Cost Matrix Construction

Using the multi-horizon success-ratio outputs described in Section IV-B, *Cord-Approx* augments proximity-based assignment with an arrival-aware success estimate. For each perceived spot $s \in \hat{\mathcal{S}}_t$, we identify its cell $k(s)$ and extract the forecast corresponding to the estimated arrival time of driver d. Specifically, for each candidate pair (d, s), we map the travel time $\tau(d, s)$ to a horizon index using Eq. (9), and define the arrival-aware success-ratio estimate as $p^{\mathrm{arr}}(d, s) = \hat{p}_{k(s),t}^{(h(d,s))}$. We then define an effective distance,

$$\tilde{\tau}(d, s) = \frac{\tau(d, s)}{p^{\mathrm{arr}}(d, s)},$$

so that candidates with a lower expected success ratio at arrival receive a larger cost. This biases matching toward spots that are both nearby and more likely to yield a successful parking outcome upon arrival. In the robust *Cord-Approx* dispatcher, this effective distance is used directly as the assignment cost, i.e., $\mathbf{M}_t(d, s) = \tilde{\tau}(d, s)$. Following our prior *Cord-Approx* design, the predicted success ratio is used to elongate travel time for spots that are less likely to remain available, and the Hungarian algorithm then minimizes this current-step effective-distance cost.

D. Assignment Algorithm and Phantom Spot Handling

The runtime dispatcher executes at every simulation step t as described in Algorithm 1. It combines the replacement-based sensor view, the arrival-aware success ratio, and the Hungarian algorithm to produce assignments for active participants. Crucially, the algorithm handles false vacancies (phantom spots). If a participant d travels to the assigned target s and finds that s is physically occupied ($s \notin \mathcal{S}_t$), then the assignment is treated as a phantom-spot encounter and is marked as failed. The participant increments a phantom-spot encounter counter (used in the evaluation metrics) and immediately enters a rerouting state, as defined in Section V. They remain active in the system and are considered for a new assignment in the subsequent time step $t + 1$. This effectively models the time penalty of misinformation. The replacement-based noise model enforces a conservation constraint: sensing errors cannot create additional parking spots. Algorithm 1 is executed once per simulation step. The predicate d.arrived() is evaluated by the simulator when a participant reaches its current target at a later step; if that target is then found occupied, the participant enters the rerouting state and is reconsidered in the next dispatch cycle.

Algorithm 1 Robust Coordinated Approximate (*Cord-Approx*) Assignment Dispatcher

Input: D_t — participants active at step t; $\mathcal{S}_t$ — ground-truth available spots; ρ, ϕ — sensor coverage and false vacancy rate; $\Delta t_f, H$ — forecasting parameters
Output: $\mathcal{A}_t^{partc}$ — assignments
Phase 1: Sensor Noise (Replacement Model)
$\mathcal{S}_{valid} \leftarrow \{s \in \mathcal{S}_t \mid \mathrm{random}() < \rho\}$
$\mathcal{G} \leftarrow \{s \notin \mathcal{S}_t \mid \mathrm{random}() < \phi\}$
$\hat{\mathcal{S}}_t \leftarrow \mathrm{ReplaceSubset}(\mathcal{S}_{valid}, \mathcal{G})$
▷ Ensures $|\hat{\mathcal{S}}_t| = |\mathcal{S}_{valid}| \le |\mathcal{S}_t|$
if $\hat{\mathcal{S}}_t$ is empty **then**
  return $\emptyset$
end if
Phase 2: Cost Matrix & Optimization
Initialize $\mathbf{M}_t$ of size $|D_t| \times |\hat{\mathcal{S}}_t|$
for $d \in D_t$ **do**
  for $s \in \hat{\mathcal{S}}_t$ **do**
    $\hat{\mathbf{p}}_{k(s),t} \leftarrow \mathbf{PredictMultiHorizon}(k(s), t)$
    $h(d, s) \leftarrow \mathbf{ArrivalHorizon}(\tau(d, s), \Delta t_f, H)$
    $p^{\mathrm{arr}}(d, s) \leftarrow \hat{p}_{k(s),t}^{(h(d,s))}$
    $\mathbf{M}_t(d, s) \leftarrow \frac{\tau(d,s)}{p^{\mathrm{arr}}(d,s)}$
  end for
end for
$\mathcal{A}_t^{partc} \leftarrow \mathbf{HungarianAssign}(\mathbf{M}_t)$
Phase 3: Dispatch & Phantom Spot Handling
for $(d, s) \in \mathcal{A}_t^{partc}$ **do**
  d.target $\leftarrow s$
  if d.arrived() **and** $s \notin \mathcal{S}_t$ **then**
    d.phantom_count $\leftarrow d$.phantom_count $+ 1$
    d.state $\leftarrow$ Rerouting ▷ Phantom Spot Encounter
  end if
end for
return $\mathcal{A}_t^{partc}$

V. Simulation Model and Experimental Setup

To evaluate robustness, we use a city-scale simulator calibrated with Madrid, Spain, traffic-intensity data [6]. The

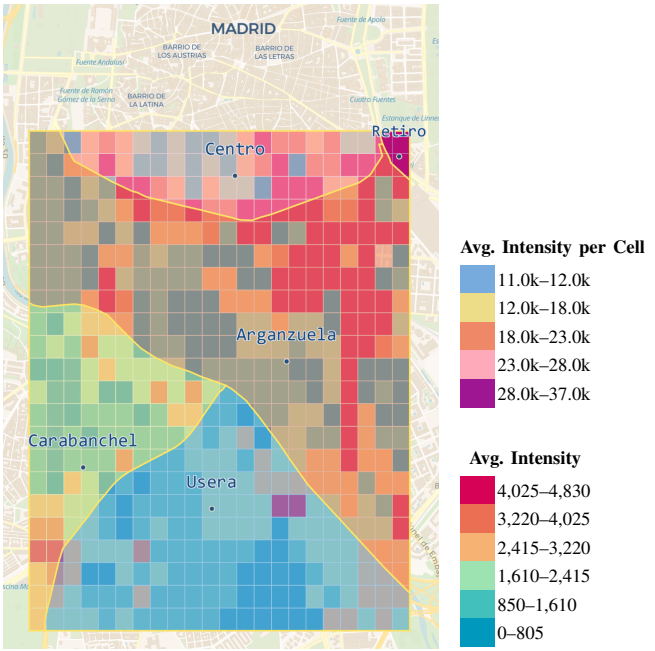

Fig. 1. Traffic-intensity heatmap of central Madrid over the 22×22 Geohash grid, averaged over peak hours (09:00–17:00). District boundaries are overlaid to show how demand varies across administrative areas. The cell-level layer colors each grid cell by its own average intensity, while the district-level layer colors each district by the mean of the cell-level average intensities of the grid cells contained in that district. The two legends report the scales for the overlaid district-level and cell-level intensity layers.

simulator tracks agents' movements, spot availability, and sensor degradation over a 24-hour horizon ($T = 1440$).

A. Simulation Environment and City Grid

The spatial domain models the street parking ecosystem of central Madrid, Spain. We overlay a geospatial grid of $N \times N$ cells ($N = 22$), where each cell corresponds to a 7-character Geohash ($\approx 150\,\mathrm{m} \times 115\,\mathrm{m}$) [21]. This mesoscopic resolution balances tractability with spatial specificity appropriate for assignment-level evaluation rather than lane-level traffic dynamics. It is therefore a modeling abstraction: the grid and Manhattan metric approximate assignment-scale travel time and do not encode road-network topology, one-way streets, congestion, or signal delays. As illustrated in Fig. 1, the average traffic intensity per cell during peak daytime hours (09:00–17:00) is shown as a heatmap overlaid with the official administrative boundaries. These boundaries partition the grid into five distinct municipal districts: **Centro**, **Retiro**, **Arganzuela**, **Carabanchel**, and **Usera**. Fig. 1 and Fig. 6 were generated with kepler.gl; the basemap uses CARTO Voyager, and the underlying map data are © OpenStreetMap contributors.

B. Parking Supply and Traffic Demand

The simulation is grounded in two primary municipal datasets [6], as follows:

• **Parking supply:** For each 7-character Geohash, we estimate the maximum on-street parking capacity by summing legally parkable street lengths while accounting for zoning, road class, and municipal restrictions (loading zones, no-stopping zones, etc.), thereby obtaining a reproducible supply ceiling. In the simulation, agents park and depart according to real-world dwell times, which dynamically updates each Geohash cell's spot availability. We adopt this upper-bound baseline because producing a complete city-wide street-space inventory remains difficult.

• **Traffic intensity:** Dynamic vehicle arrival rates derived from over 2 million sensor observations. We process 15-minute aggregated counts and disaggregate them into 1-minute bins via uniform split to generate a continuous demand profile. Following urban cruising studies, FHWA guidance [22], and our prior simulation setup, we treat roughly 10% of vehicle counts as parking-search demand, modeling 1.5% as participants and 8% as competitors, with the remaining 90.5% treated as through-traffic outside the simulation scope. Over the 24-hour horizon, this generates approximately 270,000 participants and 1.6 million competitors.

C. Agent Dynamics and Implementation Details

Agents interact within the simulation environment under a set of explicitly defined behavioral and operational constraints.

• **Participants (D_t):** Drivers coordinated by the *Cord-Approx* dispatcher. They navigate using Manhattan distance ($\tau(d, s)$) toward system-assigned spots. If a participant arrives at a false vacancy, they incur a time penalty and re-enter the pool in a rerouting state. On a phantom encounter, the participant has already spent travel time to reach the target cell; the next simulation tick, they request reassignment from their current location and continue moving toward the new target.

• **Competitors (C_t):** Background traffic relying on greedy visual search. They operate entirely independent of the sensor network and claim physical spots within a visual observability radius of $R = 1$ (observing the current and four adjacent cells).

All searching agents who cannot find parking within $t_{\max} = 30$ minutes exit the on-street search process. They must either switch to off-street/private parking or abandon the attempt, both of which are out of scope here. Search time counts all elapsed simulation minutes from agent entry until parking success or timeout, including travel and any stationary waiting ticks. VKT accumulates only when the vehicle moves between cells (Manhattan steps), converted to kilometers using the grid cell dimensions. If a participant has no assigned target in a given tick, the vehicle remains stationary for that tick, so search time increases while VKT does not. Each configuration is run with 5 random seeds; we report mean values.

D. Experimental Design and Noise Injection

To quantify robustness under increasing information degradation, we evaluate the system under varying levels of sensing error using the replacement-based sensor noise model:

• **Coverage Sweep (ρ):** $\rho \in \{0.9, 0.8, 0.7, 0.6\}$ while holding $\phi = 0.0$. This isolates the impact of sensor coverage throughout the city (False Negatives).

• **Noise Sweep (ϕ):** $\phi \in \{0.05, 0.08, 0.15, 0.20\}$ evaluated

TABLE II
MULTI-HORIZON FORECASTING PERFORMANCE ACROSS MODELS. THE PROPOSED PATCHTST-LITE (OURS) ACHIEVES THE LOWEST MEAN ABSOLUTE ERROR (MAE) AND MEAN ABSOLUTE PERCENTAGE ERROR (MAPE) ACROSS ALL EVALUATED FORECASTING HORIZONS ($h \in \{1, 2, 3\}$).

	MAE			MAPE (%)		
Method	h=1	h=2	h=3	h=1	h=2	h=3
Ridge Regression	0.158	0.172	0.174	18.19%	21.09%	21.36%
CNN	0.161	0.167	0.169	16.32%	16.97%	17.29%
PatchMLP	0.156	0.159	0.161	14.79%	15.61%	16.33%
PatchTST-lite	**0.143**	**0.148**	**0.151**	**11.6%**	**11.8%**	**12.4%**

across the coverage settings above. This measures the destructive impact of false positives.

VI. RESULTS

Building on the ideal-sensing baseline of [2], we now evaluate the upgraded *Cord-Approx* pipeline when the coordinator acts on $\hat{\mathcal{S}}_t \neq \mathcal{S}_t$ under replacement-based sensing noise. The results below first validate the multi-horizon predictor under full coverage and then quantify how parking success ratio, search time, and VKT degrade across (ρ, ϕ) relative to visually searching competitors. Unless otherwise stated, the 77.54% prior *Cord-Approx* finding and the 85.32% *Cord-Oracle* upper bound are taken from the perfect sensing evaluation reported in [2] and are used here as reference points for comparison.

A. Prediction Module Performance

Our prior work used prediction to estimate future parking success for assignment; here, the predictive layer is upgraded to a multi-horizon model that forecasts cell-level parking success ratios and provides an arrival-aware success estimate for the assignment cost. We first validate this model under full coverage before using it in the robustness experiments under Coverage Sweep (ρ) and Noise Sweep (ϕ).

To benchmark the predictive layer, we compare PatchTST-lite against Ridge Regression, a temporal CNN, and PatchMLP. Table II reports the Mean Absolute Error (MAE) and Mean Absolute Percentage Error (MAPE) for three look-ahead steps ($h \in \{1, 2, 3\}$), where h denotes the prediction of the future parking success ratio at time $t+h$ given the context window up to time t. All forecasting models are trained over the full set of grid cells; Fig. 2 shows one representative cell (cell 240) for qualitative illustration.

Table II shows that PatchTST-lite is the most accurate model at every horizon. Relative to the strongest baseline, PatchMLP, it reduces MAE from 0.156 to 0.143 at h=1 (an 8.3% reduction) and from 0.161 to 0.151 at h=3 (a 6.2% reduction). The gain is even more pronounced on MAPE: 11.6% vs. 14.79% at h=1 (21.6% lower) and 12.4% vs. 16.33% at h=3 (24.1% lower). These improvements matter operationally because the arrival-aware success forecast enters the assignment cost directly: smaller horizon errors mean fewer misranked candidate spots and fewer avoidable reroutes. Figure 2 illustrates why these forecasting gains matter. PatchTST-lite tracks regime shifts more faithfully while suppressing the spurious oscillations visible in the baseline models. In the assignment pipeline, this added stability matters because forecast errors are not merely descriptive; they directly affect which spots are prioritized for dispatch.

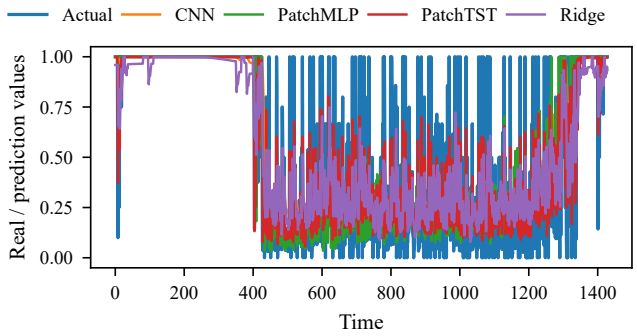

Fig. 2. Example time series for one grid cell (cell 240), shown for qualitative illustration, although all models are trained on the full set of cells. The ground-truth parking success ratio series and the predicted multi-horizon trajectories are plotted over simulation time to compare Ridge regression, CNN, PatchMLP, and PatchTST-lite.

B. System Performance and Diurnal Dynamics

We next evaluate system-level performance over the full 24-hour horizon and during peak hours (09:00–17:00). The question is whether the upgraded *Cord-Approx* pipeline preserves a coordinated advantage over visually searching competitors under reduced coverage in the evaluated setting and, under perfect sensing, how much of the prior gap to the *Cord-Oracle* upper bound can be recovered. Figure 3 compares hourly parking success ratio for participants across sensing-coverage levels and against the competitors baseline. Only one competitors curve is shown because competitors do not use the sensing system and are therefore not directly parameterized by ρ; their curve serves as the visual-search baseline throughout the coverage sweep. During peak hours (09:00–17:00), the competitors baseline remains far below the coordinated participants curves, while the five participants curves stay relatively close to one another. Under perfect sensing (PS: ρ=1, ϕ=0), the upgraded predictor increases the *Cord-Approx* parking success ratio to 82.2%, up from 77.54% in [2], thereby moving closer to the *Cord-Oracle* upper bound of 85.32%.

The peak-hour zoom makes the bounded degradation under sensor coverage much easier to read. In that window, the largest separation between the best-coverage and worst-coverage participants curves, i.e., between $\rho = 1.0$ and $\rho = 0.6$, is $\Delta_{\max} = 0.100$ at hour 11, while the smallest is $\Delta_{\min} = 0.047$ at hour 9. Since the y-axis reports parking success ratio on the fractional interval $[0, 1]$, these gaps correspond to 10.0 and 4.7 percentage points, respectively. This is important for interpretation: even though the curves are visually separated, the peak-hour degradation induced by reducing coverage from $\rho = 1.0$ to $\rho = 0.6$ remains below

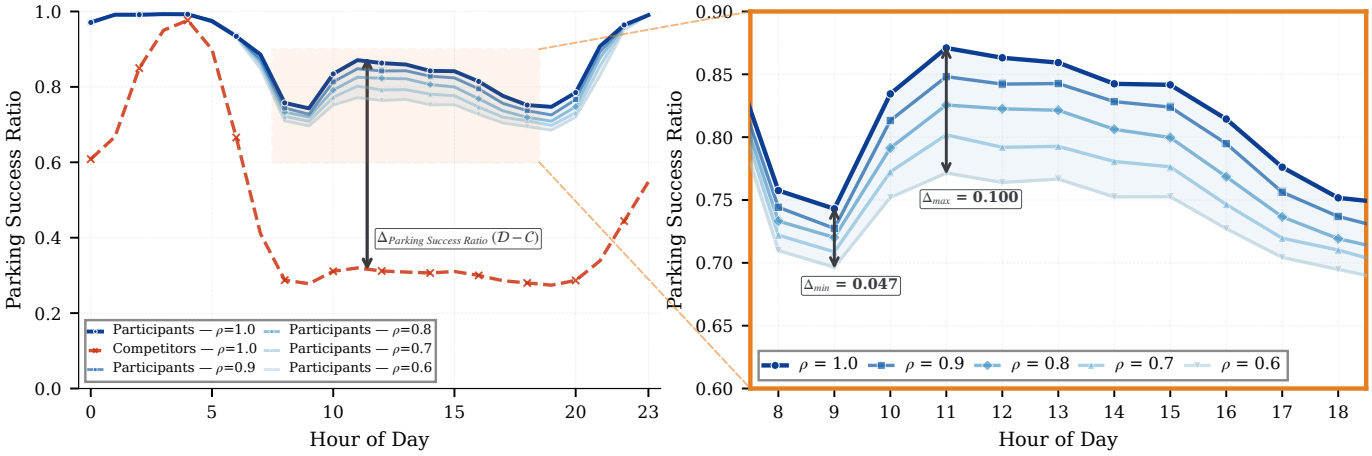

Fig. 3. Hourly parking success ratio over a 24-hour horizon for participants under sensor-coverage levels $\rho \in \{0.6, 0.7, 0.8, 0.9, 1.0\}$ with $\phi = 0$, compared with the competitors baseline. The competitors curve is shown once because competitors search visually and do not rely on the sensing system. The right panel zooms into peak hours (09:00–17:00) and highlights the spread between the best-coverage and worst-coverage participants curves, i.e., between $\rho = 1.0$ and $\rho = 0.6$. The smallest observed spread in this window is $\Delta_{\min} = 0.047$ at hour 9 (4.7 percentage points), and the largest is $\Delta_{\max} = 0.100$ at hour 11 (10.0 percentage points), since the parking success ratio is plotted on a fractional $[0, 1]$ scale.

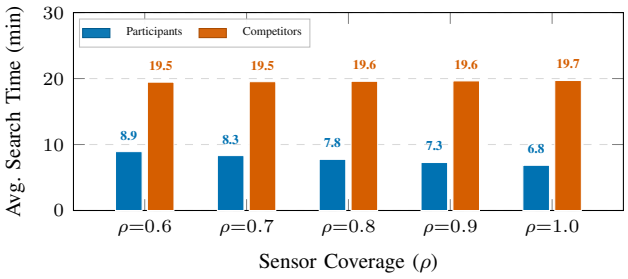

(a) Peak hours (09:00–17:00)

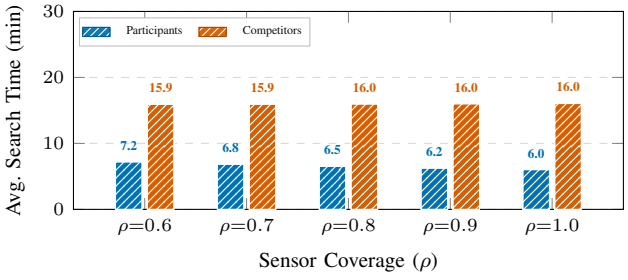

(b) Off-peak hours

Fig. 4. Average search time versus sensor coverage ρ for participants and competitors. Panel (a) reports peak hours (09:00–17:00) and panel (b) off-peak hours (hatched). Across the peak-hour sensor coverage sweep, participants remain 54.1%–65.2% faster than competitors, and across the off-peak sweep they remain 54.8%–62.5% faster, showing that sensing degradation raises participants search time but does not erase the coordinated time advantage.

0.10 in fractional units throughout the annotated window. In other words, ρ reduces participants performance, but the resulting peak-hour loss is bounded and does not erase the coordinated advantage over visually searching competitors.

The temporal efficiency of coordination is further decomposed in Fig. 4, which profiles average search time across the ρ sweep. During peak hours (panel a), competitors remain near 19.5–19.7 minutes regardless of the sensing state because they do not use the system. By contrast, guided participants average 6.85 minutes under perfect sensor coverage ($\rho = 1.0$), versus 19.70 minutes for competitors, i.e., a 65.2% reduction in search time. Even when coverage falls to $\rho = 0.6$, participants rise only to 8.92 minutes, which is still 54.1% lower than the 19.45-minute competitor baseline. Thus, losing 40% of detections raises participants' peak-hour search time by 30.2% relative to perfect coverage (6.85 to 8.92 minutes), but leaves a double-digit minute advantage intact. The same pattern persists off-peak (panel b), where participants remain 54.8%–62.5% faster than competitors across the coverage sweep.

C. Robustness to Information Omission (ρ)

To isolate omissions, we examine Table III at $\phi = 0.00$. As coverage drops from $\rho = 0.9$ to $\rho = 0.6$, participants success falls from 80.4% to 74.2% (a drop of 6.2 percentage points, or 7.7% relative), search time rises from 7.27 to 8.92 minutes (+22.7%), and VKT falls from 0.53 to 0.37 km (-30.2%). This opposite movement is the key omission effect: lower sensor coverage hides real supply from the coordinator, so physical vacancies remain unclaimed and accumulate closer to incoming drivers, which shortens travel distance, yet the incomplete perceived map also produces more missed opportunities and failed assignments, which increases elapsed search time. Operationally, omission therefore hurts guidance quality more in time than in distance. Even in the lowest-coverage regime, participants still retain a +43.1 percentage-point success advantage and a 54.1% search-time reduction over competitors. The bold rows in Table III identify a coverage-dependent sub-10-minute operating regime: the threshold is maintained up to $\phi = 0.05$ at $\rho = 0.6$, $\phi = 0.08$ at $\rho = 0.7$ and $\rho = 0.8$, and $\phi = 0.15$ at $\rho = 0.9$.

D. The Penalty of Misinformation (ϕ)

While omission hides real parking supply, false vacancies actively mislead the dispatcher. At $\rho = 0.9$, increasing ϕ

TABLE III

PEAK-HOUR (09:00–17:00) PERFORMANCE UNDER VARYING SENSOR COVERAGE (ρ) AND FALSE VACANCY RATE (ϕ). BOLD ENTRIES MARK, FOR EACH COVERAGE LEVEL ρ, THE LARGEST TESTED ϕ FOR WHICH AVERAGE PARTICIPANTS SEARCH TIME REMAINS BELOW 10 MINUTES; THE CORRESPONDING PARTICIPANTS PARKING SUCCESS RATIO IS BOLDED ON THE SAME ROW FOR REFERENCE. THE TABLE SHOWS TWO COMPLEMENTARY ROBUSTNESS PATTERNS: SENSOR COVERAGE MAINLY TRADES OFF TIME AGAINST DISTANCE, WHEREAS FALSE VACANCY RATE PRIMARILY INCREASES SEARCH TIME THROUGH REROUTING.

ρ	ϕ	Success Ratio (%)		Search Time (min)		VKT (km)	
		Participants	Competitors	Participants	Competitors	Participants	Competitors
0.6	0.00	74.2	31.1	8.92	19.45	0.37	2.32
0.6	0.05	**72.6**	30.7	**9.42**	19.57	0.33	1.87
0.6	0.08	70.4	30.2	10.15	19.71	0.37	1.88
0.6	0.15	64.4	30.8	11.83	19.51	0.35	1.86
0.6	0.20	60.9	31.2	12.99	19.40	0.34	1.85
0.7	0.00	76.4	30.9	8.32	19.52	0.43	2.33
0.7	0.05	75.2	30.6	8.72	19.59	0.37	1.87
0.7	0.08	**73.4**	30.2	**9.22**	19.72	0.40	1.88
0.7	0.15	66.9	30.7	11.10	19.54	0.41	1.86
0.7	0.20	63.4	30.9	12.20	19.51	0.40	1.86
0.8	0.00	78.5	30.7	7.75	19.58	0.48	2.33
0.8	0.05	77.6	30.4	8.03	19.66	0.41	1.87
0.8	0.08	**76.2**	30.1	**8.49**	19.76	0.44	1.88
0.8	0.15	69.6	30.3	10.41	19.63	0.46	1.87
0.8	0.20	66.1	30.6	11.29	19.58	0.46	1.87
0.9	0.00	80.4	30.5	7.27	19.64	0.53	2.34
0.9	0.05	79.8	30.3	7.50	19.70	0.45	1.88
0.9	0.08	78.9	30.0	7.81	19.80	0.47	1.88
0.9	0.15	**72.9**	29.9	**9.51**	19.76	0.53	1.88
0.9	0.20	68.5	30.4	10.57	19.61	0.51	1.87
1.0	0.00	82.2	30.3	6.85	19.70	0.58	2.35

from 0.00 to 0.15 raises participants search time from 7.27 to 9.51 minutes (+30.8%) and reduces participants success from 80.4% to 72.9% (-7.5 pp), while the participants–competitors success ratio gap shrinks from +49.94 to +43.00 percentage points. Because VKT changes much less over the same rows, the main cost of misinformation is temporal rather than spatial: drivers lose time resolving bad assignments rather than driving much farther. Table III also shows the secondary redistribution effect on competitors. At $\rho = 0.6$, competitor VKT falls from 2.32 to 1.86 km (-19.8%) between $\phi = 0.00$ and $\phi = 0.15$ because degraded coordinated assignments leave more real spots available to background traffic. Thus false vacancies do not help competitors directly; they weaken the coordinated system enough to free parking spots.

E. City-Scale Operational Boundaries

We define the operational boundary using the participants–competitors parking success-ratio gap in Fig. 5(b) and examine how that gap evolves across ϕ. In this subsection, the gap is the primary boundary criterion; the 10-minute search-time threshold from Table III is used only as a complementary operational interpretation. Panel (a) shows that participants' peak-hour parking success ratio declines from the perfect sensing reference of 82.22% at $(\rho, \phi) = (1.0, 0)$ to 60.86% at the noisiest evaluated point $(0.6, 0.20)$, a drop of 21.36 percentage points. Panel (b) reports the more decision-relevant quantity: the advantage over competitors. Under perfect sensing, that advantage is +51.90 percentage points; along the $\phi = 0.15$ row, it remains between +33.59 and +43.00 percentage points across all tested coverages, but it falls further to +29.63–+38.07 percentage points at $\phi = 0.20$. This is why $\phi \approx 0.15$

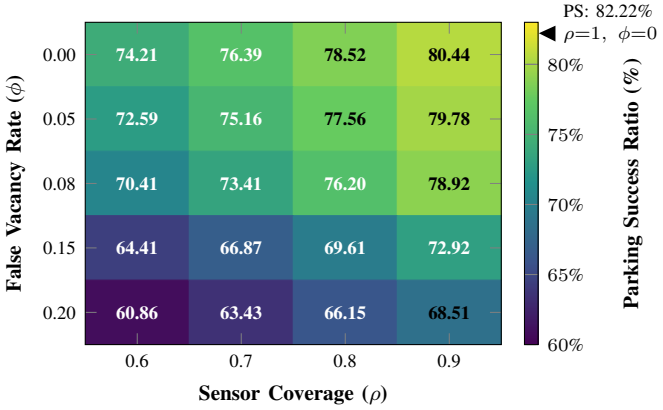

(a) Average parking success ratio of participants during peak hours (09:00–17:00).

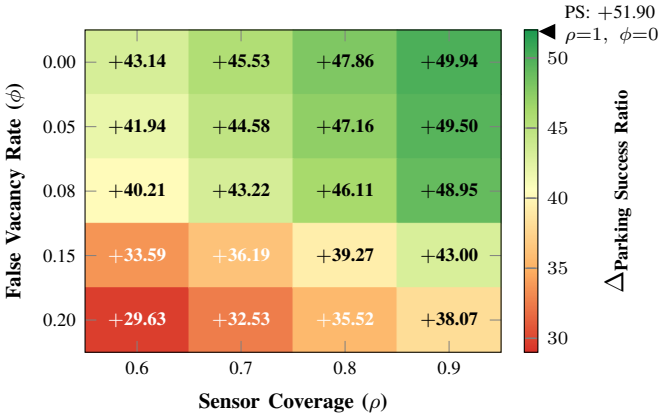

(b) Performance advantage of participants over competitors (pp).

Fig. 5. Peak-hour robustness across the $\rho \times \phi$ sweep. Panel (a) reports participants' parking success ratio, and panel (b) reports the participants advantage over competitors in percentage points. The external Perfect Sensing (PS) marker (ρ=1, ϕ=0) denotes the perfect sensor coverage reference excluded from the noisy grid. Performance degrades as false vacancies increase, but the coordinated advantage remains positive throughout the evaluated range.

is the practical boundary in the evaluated setting under the parking success gap criterion: it is the last tested regime that still preserves a large success advantage across all evaluated coverage levels. Importantly, this is not a collapse point: even at $(\rho, \phi) = (0.6, 0.20)$, participants still retain a +29.7 percentage-point success advantage and a 33.0% search-time reduction (12.99 vs. 19.40 minutes) over competitors.

F. District-Level Robustness Under Noise

To expose the spatial heterogeneity hidden by city-wide averages, we examine a representative noisy operating point, $(\rho, \phi) = (0.7, 0.08)$, in district-aggregated form and compare district-level mean search times in Fig. 6. Each district value is obtained by averaging the cell-level mean search times over the grid cells assigned to that district, so the district summaries remain directly comparable with the grid-based analysis used throughout the paper. At this operating point, participants outperform competitors in all five districts, but the magnitude of the gain varies substantially. Search time is

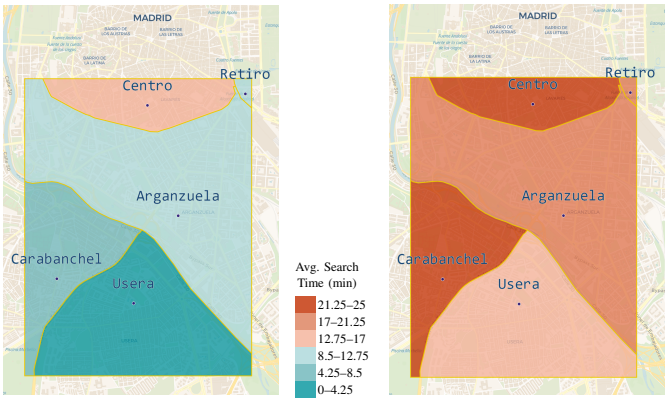

(a) Avg. search time for participants.

(b) Avg. search time for competitors.

Fig. 6. District-level average search time for the representative noisy setting $(\rho, \phi) = (0.7, 0.08)$. Panel (a) reports participants and panel (b) reports competitors using a shared legend. Each district value is obtained by averaging the cell-level mean search times over the grid cells assigned to that district. This shows that the coordinated advantage survives in every district, although its magnitude varies substantially across the map.

reduced by 75.1% in Carabanchel (5.51 vs. 22.11 minutes) and 74.5% in Usera (3.53 vs. 13.83 minutes), while the reduction is smaller but still substantial in Centro at 32.3% (16.91 vs. 24.97 minutes). Arganzuela and Retiro lie between these extremes, with 54.5% and 53.1% reductions, respectively. The district context helps explain why these reductions differ in magnitude. Centro is the historic, tourism- and service-intensive core of Madrid, so higher curbside competition, visitor activity, and dense central demand are consistent with its remaining the hardest district even under coordination. Retiro also combines strong cultural and visitor pull with a central location, which is compatible with intermediate rather than extreme gains. By contrast, Arganzuela functions as the southern extension of the historic center and includes large redeveloped and riverfront areas around Madrid Río, while Carabanchel and Usera belong to Madrid's southern districts; in this aggregated view, these less central districts exhibit lower mean search times and larger relative gains. The key point is therefore not that coordinated assignment works only in easier residential areas, but that it remains beneficial across distinct urban contexts, from the tourism-intensive center to the southern districts with lower baseline pressure.

G. Key Takeaways

We summarize the main empirical observations from the robustness sweeps over sensor coverage (ρ) and false vacancy rate (ϕ) in mixed traffic (participants vs. competitors):

- **Improved perfect-sensing baseline:** Under perfect sensing ($\rho = 1$, $\phi = 0$), the upgraded predictor raises the participants parking success ratio to 82.2%, a gain of 4.66 percentage points (pp) over the 77.54% value reported in our prior work and a reduction of the gap to the 85.32% *Cord-Oracle* upper bound from 7.78 to 3.12 pp.
- **Robustness to coverage loss:** Even at $\rho = 0.6$ with $\phi = 0$, participants achieve a 74.2% parking success ratio versus 31.1% for competitors, corresponding to both a +43.1 percentage-point absolute advantage and a 2.39× relative ratio, while reducing average search time by 54.1% (8.92 vs. 19.45 minutes).
- **Temporal penalty of false vacancies:** At $\rho = 0.9$, increasing ϕ from 0 to 0.15 raises participants search time by 30.8% (7.27 to 9.51 minutes) and shrinks the participants–competitors success-ratio gap from +49.94 to +43.00 percentage points, showing that misinformation erodes coordinated gains before it eliminates them.
- **Operational threshold in the evaluated setting:** Using the participants–competitors parking success ratio gap as the operational criterion, coordinated benefits remain strongly positive up to $\phi \approx 0.15$; along that row, the gap still ranges from +33.59 to +43.00 pp across the tested coverage levels, before weakening further at $\phi = 0.20$.
- **District-level heterogeneity:** At the representative noisy setting $(\rho, \phi) = (0.7, 0.08)$, the average search-time reduction ranges from 32.3% in Centro to 75.1% in Carabanchel and 74.5% in Usera, showing that robustness persists across both central and southern districts in the aggregated view.

VII. Conclusion

We quantified the robustness of coordinated on-street parking assignment under imperfect sensing, partial sensor coverage (ρ), and false vacancy rate (ϕ), using a replacement-based noise model that preserves physical capacity. Under perfect sensing ($\rho = 1$, $\phi = 0$), the PatchTST-lite predictor raises the parking success ratio to 82.2%, a gain of 4.66 pp (6.0% relative) over the 77.54% baseline reported in [2], and reduces the gap to the 85.32% *Cord-Oracle* upper bound from 7.78 to 3.12 pp. Under degraded coverage, even when 40% of detections are missing ($\rho = 0.6$, $\phi = 0$), participants achieve a 74.2% parking success ratio versus 31.1% for competitors (+43.1 percentage points, or 2.39× in relative ratio) and reduce search time by 54.1% (8.92 vs. 19.45 minutes).

This advantage is not limited to city-wide averages: in the district-level analysis for a representative noisy setting, participants also achieve lower mean search time than competitors in all five Madrid districts, including the most crowded district, Centro. As ϕ increases, the coordinated advantage erodes mainly through rerouting rather than longer travel, with the participants–competitors parking success-ratio analysis indicating an empirical operating region up to roughly $\phi \approx 0.15$ in the evaluated setting. The main contribution of the extension is therefore not only improved prediction under perfect sensing, but also a physically consistent robustness framework for determining when coordinated assignment remains operational under degraded sensing. Future work will study transfer to other cities by recalibrating demand, parking supply, adoption mix, and travel-time metrics, and will incorporate road-network distances, one-way streets, signal delays, congestion and emissions impacts, correlated sensor failures, user response to repeated misguidance, field validation, and prototype implementation.

REFERENCES

[1] D. Shoup, "Cruising for parking," *Access Magazine*, vol. 1, no. 30, pp. 16–23, 2007.

[2] 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*, SIGSPATIAL '25, pp. 378–381, Association for Computing Machinery, 2025.

[3] C. Zimmerman, R. Klein, J. Schroeder, K. Turnbull, K. Balke, M. Burris, E. Saunoi-Sandgren, E. Martin, S. Shaheen, C. Rodier, *et al.*, "San francisco urban partnership agreement: national evaluation report.," tech. rep., United States. Department of Transportation. Intelligent Transportation, 2014.

[4] S. Mathur, T. Jin, N. Kasturirangan, J. Chandrasekaran, W. Xue, M. Gruteser, and W. Trappe, "ParkNet: Drive-by sensing of road-side parking statistics," in *Proceedings of the 8th International Conference on Mobile Systems, Applications, and Services (MobiSys '10)*, pp. 123–136, Association for Computing Machinery, 2010.

[5] B. Xu, O. Wolfson, J. Yang, L. Stenneth, P. S. Yu, and P. C. Nelson, "Real-time street parking availability estimation," in *2013 IEEE 14th International Conference on Mobile Data Management*, vol. 1, pp. 16–25, IEEE, 2013.

[6] Madrid City Council, "Trafico. Historico de datos del trafico desde 2013." https://datos.madrid.es/dataset/208627-0-transporte-ptomedida-historico, 2024. Accessed: 2024-04-18.

[7] M. Fosgerau and A. de Palma, "The dynamics of urban traffic congestion and the price of parking," *Journal of Public Economics*, vol. 105, pp. 106–115, 2013.

[8] G. Dalla Chiara, K. F. Krutein, A. Ranjbari, and A. Goodchild, "Providing curb availability information to delivery drivers reduces cruising for parking," *Scientific Reports*, vol. 12, no. 1, p. 19355, 2022.

[9] S. Nawaz, C. Efstratiou, and C. Mascolo, "ParkSense: A smartphone-based sensing system for on-street parking," in *Proceedings of the 19th Annual International Conference on Mobile Computing and Networking (MobiCom '13)*, MobiCom '13, pp. 75–86, Association for Computing Machinery, 2013.

[10] S. Inam, A. Mahmood, S. Khatoon, M. Alshamari, and N. Nawaz, "Multisource data integration and comparative analysis of machine learning models for on-street parking prediction," *Sustainability*, vol. 14, no. 12, p. 7317, 2022.

[11] 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.

[12] S. Zambanini, A.-M. Loghin, N. Pfeifer, E. M. Soley, and R. Sablatnig, "Detection of parking cars in stereo satellite images," *Remote Sensing*, vol. 12, no. 13, p. 2170, 2020.

[13] C. Roman, R. Liao, P. Ball, S. Ou, and M. de Heaver, "Detecting on-street parking spaces in smart cities: Performance evaluation of fixed and mobile sensing systems," *IEEE Transactions on Intelligent Transportation Systems*, vol. 19, no. 7, pp. 2234–2245, 2018.

[14] Y. Geng and C. G. Cassandras, "New "smart parking" system based on resource allocation and reservations," *IEEE Transactions on Intelligent Transportation Systems*, vol. 14, no. 3, pp. 1129–1139, 2013.

[15] B. Q. Tan, S. X. Xu, M. Thürer, K. Kang, Z. Zhao, and M. Li, "Booking versus search-based parking strategy: A game-theoretic methodology," *Research in Transportation Economics*, vol. 104, p. 101416, 2024.

[16] A. O. Kotb, Y.-C. Shen, X. Zhu, and Y. Huang, "iParker—a new smart car-parking system based on dynamic resource allocation and pricing," *IEEE Transactions on Intelligent Transportation Systems*, vol. 17, no. 9, pp. 2637–2647, 2016.

[17] H. W. Kuhn, "The hungarian method for the assignment problem," *Naval Research Logistics Quarterly*, vol. 2, no. 1–2, pp. 83–97, 1955.

[18] A. Millard-Ball, R. R. Weinberger, and R. C. Hampshire, "Is the curb 80% full or 20% empty? assessing the impacts of san francisco's parking pricing experiment," *Transportation Research Part A: Policy and Practice*, vol. 63, pp. 76–92, 2014.

[19] A. Rodríguez, B. Alonso, J. L. Moura, and L. dell'Olio, "Analysis of user behavior in urban parking under different level of information scenarios provided by smart devices or connected cars," *Travel Behaviour and Society*, vol. 37, p. 100847, 2024.

[20] T. Higuchi and K. Oguchi, "On the predictability of parking preferences," in *Proceedings of the 2021 ACM International Symposium on Wearable Computers*, ISWC '21, pp. 30–31, Association for Computing Machinery, 2021.

[21] Geohash.es, "Geohash.es — geohash encoding and decoding." https://geohash.es/. Accessed: 2025-05-18.

[22] R. Weinberger, A. Millard-Ball, T. Fabusuyi, E. Calvin, J. Blackburn, and M. Neuner, "Parking cruising analysis methodology: Final project report," fhwa-hop-23-004, Federal Highway Administration, 2023.