

UrbanHuRo: A Two-Layer Human-Robot Collaboration Framework for the Joint Optimization of Heterogeneous Urban Services

Tonmoy Dey^{1†}, Lin Jiang^{1†}, Zheng Dong², Guang Wang^{1*}

Abstract—In the vision of smart cities, technologies are being developed to enhance the efficiency of urban services and improve residents’ quality of life. However, most existing research focuses on optimizing individual services in isolation, without adequately considering the reciprocal interactions among heterogeneous urban services that could yield higher efficiency and improved resource utilization. For example, human couriers could collect traffic and air quality data along their delivery routes, while sensing robots could assist with on-demand delivery during peak hours, enhancing both sensing coverage and delivery efficiency. However, the joint optimization of different urban services is challenging due to their potentially conflicting objectives and real-time coordination in dynamic environments. In this paper, we propose UrbanHuRo, a two-layer human-robot collaboration framework for joint optimization of heterogeneous urban services, demonstrated through the examples of crowdsourced delivery and urban sensing. There are two innovative designs in UrbanHuRo, i.e., (i) a scalable distributed MapReduce-based K-Submodular maximization module for efficient order dispatch and (ii) a deep submodular reward reinforcement learning algorithm for sensing route planning. Experimental evaluations on real-world datasets from a food delivery platform demonstrate that our UrbanHuRo improves sensing coverage by 29.7% and courier income by 39.2% on average in most settings, while also significantly reducing the number of overdue orders.

I. INTRODUCTION

Rapid advancements in Artificial Intelligence (AI) and sensing technologies have enabled diverse smart city services, such as large-scale real-time order dispatch in food delivery [1] and robot vehicle (RV) cruising for urban sensing [2]. While many studies [3]–[5] focus on improving the efficiency of individual services, they often treat these services in isolation, overlooking opportunities to leverage resources across services for higher overall efficiency. For instance, food-delivery couriers could collect air-quality or traffic data along their routes, while sensing RVs could assist with food delivery during peak hours, forming a win-win collaboration, as illustrated in Fig. 1.

Although a collaborative paradigm holds great potential for improving the efficiency of heterogeneous urban services, designing effective algorithms for joint optimization remains nontrivial due to several fundamental **challenges**: (i) *Conflicting objectives and asynchronous reward feedback*. Different services often have independent and even conflicting objectives, making resource sharing difficult without compromising the performance of individual tasks. Moreover, it is




Fig. 1: Intuition of Human-Robot Collaboration

challenging to assess how a single action impacts multiple objectives when rewards are observed asynchronously. For example, the potential sensing benefit of an order dispatch decision cannot be determined in advance, as it depends on subsequent sensing actions. (ii) *Real-time coordination in dynamic environments*. Coordinating large numbers of agents in real time is inherently complex and computationally demanding, particularly in urban systems involving multiple interdependent services. (iii) *Human-robot collaboration*. The coexistence of human workers (e.g., couriers) and autonomous robots (e.g., sensing RVs) adds complexity. Robots follow instructions strictly, whereas human workers have diverse preferences and behaviors. Their satisfaction and willingness to participate should be considered in decision-making to ensure sustainability.

To address these challenges, we propose **UrbanHuRo**, a two-layer human-robot collaboration framework for jointly optimizing real-time crowdsourced delivery and urban sensing services. Our objectives are threefold: minimizing overdue orders, maximizing delivery income, and enhancing urban sensing coverage. The framework introduces **hybrid reward-value feedback** to address the first challenge, i.e., asynchronous reward signals in multi-objective optimization. Specifically, the order dispatch layer makes decisions based on both immediate delivery rewards and estimated sensing values computed by the sensing layer, representing expected future sensing returns. To address the second and third challenges, we further design two coupled components in UrbanHuRo: (i) an efficient MapReduce-based K-Submodular maximization module for order dispatch, referred to as **KSubMR**. This module leverages the submodular property of the matching function and employs the MapReduce paradigm to enable parallel, real-time coordination in large-scale dispatch problems, thereby enhancing computational efficiency. (ii) a Deep Submodular Reward Q-Network (**DSRQN**) for sensing route planning. Here, human couriers follow their preferred routes to encourage participation, while RVs are guided by a coverage-aware spatio-

[†]These authors contributed equally (co-first authors).

*Corresponding author.

¹Department of Computer Science, Florida State University.

²Department of Computer Science, Wayne State University.

temporal reward function to dynamically balance sensing and delivery tasks. To facilitate human-robot coordination, sensing gains from human agents are computed directly, whereas those from RVs are passed to the upper layer and aggregated using a submodular function.

To summarize, our main contributions are as follows:

- **Conceptually**, we study the optimization of heterogeneous urban services through human-robot collaboration. In particular, we integrate two originally independent services, food delivery performed by human couriers and urban sensing conducted by RVs, for joint optimization. Our motivation is to enable each service to *opportunistically* utilize idle resources to support the other through cooperation.
- **Technically**, we propose **UrbanHuRo**, a two-layer human-robot collaboration framework designed to jointly optimize heterogeneous urban services. The upper decision layer is a scalable MapReduce-based K-Submodular maximization module, **KSubMR**, for efficient order dispatch. The lower decision layer includes a novel deep reinforcement learning algorithm called **DSRQN**, for effective sensing route planning.
- **Experimentally**, we evaluate UrbanHuRo on a real-world dataset from a food delivery platform, containing 160K orders collected in Shanghai. Experimental results demonstrate that our UrbanHuRo improves sensing coverage by 29.7% and courier income by 39.2% on average in most settings, while also significantly reducing the number of overdue orders.

II. PROBLEM FORMULATION

Considering the sequential nature of crowdsourced delivery and urban sensing tasks, we formulate their joint optimization as a Markov Decision Process (MDP), denoted by $\mathcal{M} = (\mathcal{S}, \mathcal{A}, \mathcal{R}, \mathcal{P}, \gamma)$, where \mathcal{S} is the state space, \mathcal{A} is the action space, \mathcal{R} is the reward function, $\mathcal{P}(s_{t+1} | s_t, a_t)$ specifies the state transition dynamics, and $\gamma \in [0, 1)$ is the discount factor. In our setting, an entire day is divided into multiple time slots. At each slot, agents interact with the environment following a policy $\pi(a_t | s_t)$ that selects the joint action based on the current state.

- **Agent**: We model human couriers and sensing RVs as agents. When not distinguishing the two types, we use d^i to denote the i -th agent. Specifically, the i -th courier is c^i and the i -th RV is rv^i . The total fleet sizes of couriers and RVs are N^c and N^{rv} , respectively, and are assumed constant throughout the day. In contrast, the numbers of available couriers and RVs vary over time and are denoted at time step t by n_t^c and n_t^{rv} .
- **State** \mathcal{S} : Each agent has a state $s_t^i \in \mathcal{S}$, defined as the concatenation of three components: agent-specific features as_t^i , order-related features os_t^i , and sensing-related features ss_t^i , i.e., $s_t^i = [as_t^i, os_t^i, ss_t^i]$. The agent-specific component is $as_t^i = [t, g_t^i, avail_t^i, type^i]$, where t is the current time step, g_t^i is the region (grid) ID where agent i is located, $avail_t^i$ indicates whether the agent is available, and $type^i$ specifies whether the agent is a

human courier or an RV. The order-related component $os_t^i = [o_{p,t}^i, o_{d,t}^i, o_{ddl,t}^i, o_{fee,t}^i]$ includes the pickup and drop-off regions, the promised delivery deadline, and the delivery fee. Here, os_t^i corresponds to the order(s) currently assigned to agent i (denoted by \mathcal{O}_t^i), with a capacity constraint $|\mathcal{O}_t^i| \leq c_d$. The sensing-related component is $ss_t^i = [g_t^{i,last}, neigh_t^i, prop(g_t^i)]$, where $g_t^{i,last}$ denotes the last region where agent i performed sensing, $neigh_t^i$ is the number of adjacent regions visited by agent i in the past hour, and $prop(g_t^i)$ encodes region properties such as region type and accessibility.

- **Action** \mathcal{A} : We consider a joint action $a_t = (a_t^o, a_t^s) \in \mathcal{A}$ in our two-layer decision framework, where $\mathcal{A} = \mathcal{A}^o \times \mathcal{A}^s$. In the dispatch layer, a_t^o assigns each available order $o_t^j \in \mathcal{O}_t$ to an agent d^i (a courier or an RV). In the sensing layer, we make routing decisions only for RVs: $a_t^s = \{a_t^{s,i}\}_{i=1}^{n_t^{rv}}$, where each $a_t^{s,i}$ moves an rv^i from its current region to one of the eight neighboring regions (8-connected directions). Couriers do not take routing actions in the sensing layer and are assumed to travel according to their own preferences (e.g., fast and high-reward routes), whereas RVs strictly follow a_t^s .
- **Reward** \mathcal{R} : At the end of time slot t , we define the reward received by each agent d^i during this slot t as $r_t^i = (r_{d,t}^i, r_{s,t}^i)$, where $r_{d,t}^i$ and $r_{s,t}^i$ denote the delivery and sensing rewards, respectively. Here, $r_{d,t}^i$ aggregates the delivery rewards from all orders completed by agent d^i during time slot t , and the definition is shown as:

$$r_{d,t}^i = \sum_{o^j \in \mathcal{O}_t^{i,comp}} o_{fee,t}^j \times \beta^{\Delta t_{o^j}}, \quad (1)$$

where $\mathcal{O}_t^{i,comp}$ is the set of orders completed by agent d^i during time slot t , $o_{fee,t}^j$ is the delivery fee of order o^j , Δt_{o^j} is its delivery delay, and $\beta \in (0, 1)$ is a late-delivery discount factor. If $\Delta t_{o^j} = 0$, the full fee is obtained; otherwise, it decays exponentially with Δt_{o^j} . The sensing reward $r_{s,t}^i$ is defined as:

$$r_{s,t}^i = r_{g_t^i, d^i}^{reg} + r_{g_t^i, d^i}^{nbr} - r_{g_t^i, d^i}^{pen}. \quad (2)$$

Here, g_t^i is the region visited by agent d^i at time slot t . The regional term $r_{g_t^i, d^i}^{reg}$ measures the sensing value at g_t^i and is set to zero if g_t^i has been visited within the past hour to avoid redundancy. The neighboring term $r_{g_t^i, d^i}^{nbr}$ encourages broader coverage by rewarding visits to regions adjacent to unsensed regions, computed in the same way as $r_{g_t^i, d^i}^{reg}$. The penalty term $r_{g_t^i, d^i}^{pen}$ imposes a large cost when an order assigned to d^i misses its deadline, discouraging risky sensing actions that may delay deliveries.

- **Transition** \mathcal{P} : A mapping $\mathcal{S} \times \mathcal{A} \times \mathcal{S} \rightarrow [0, 1]$, where $\mathcal{P}(s_{t+1} | s_t, a_t)$ denotes the probability of transitioning from s_t to s_{t+1} under the action a_t . In UrbanHuRo, \mathcal{P} is driven by within-step updates of agent locations and order statuses in the Environment and Simulator modules: routing actions a_t^s move RVs across regions, while the dispatch action a_t^o assigns newly arrived orders to both couriers and RVs.

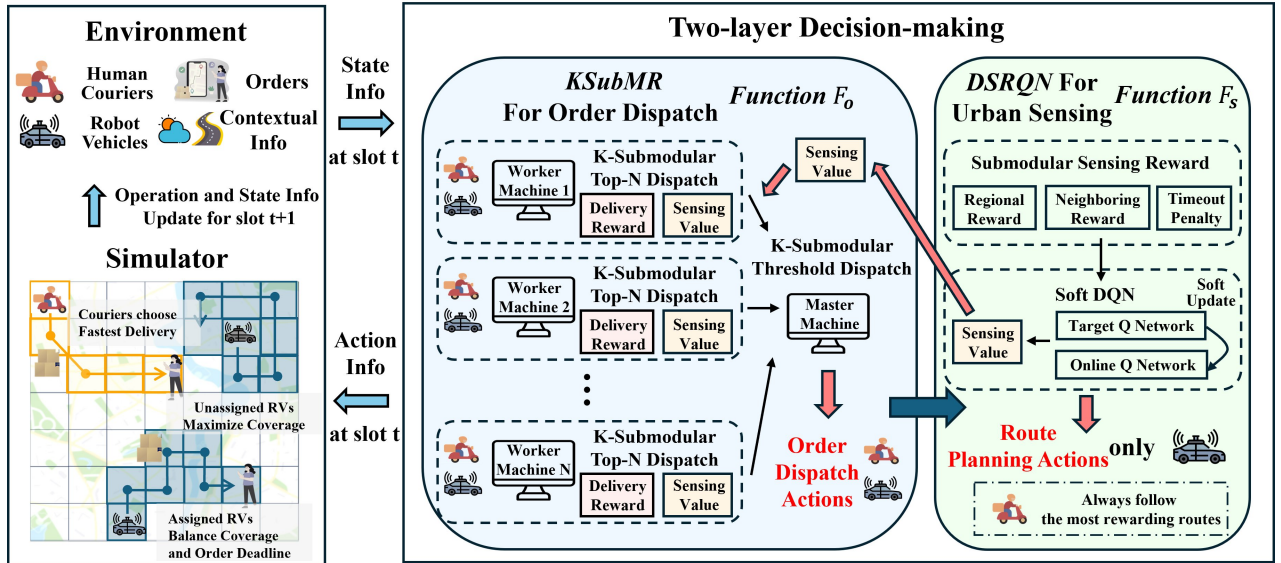


Fig. 2: An Overview of the Proposed UrbanHuRo Framework

- **Policy** π : In an MDP, a policy π maps states to actions. UrbanHuRo employs two policies in a two-layer decision framework and updates them at each time step: the order dispatch policy F_o , which generates the dispatch action a_t^o for all agents, and the urban sensing policy F_s , which produces routing actions a_t^s for RVs.

III. DESIGN OF URBANHURO

To solve the above MDP, we propose a two-layer human-robot collaboration framework, **UrbanHuRo**. As shown in Fig. 2, UrbanHuRo comprises the *Environment*, the *Simulator*, and a *Decision-making* module with two coupled layers, order dispatch and route planning. In the **order dispatch** layer, the system assigns orders to both human couriers and RVs. In the **route planning** layer, couriers follow their preferences (i.e., profit-maximizing routes), whereas RVs strictly follow system-generated actions for sensing.

A. Environment and Simulator

In the *Environment* module, we model couriers, RVs, and orders, together with contextual factors such as weather and the road network. Each order includes pickup and drop-off regions, a delivery fee, and a deadline. Both couriers and RVs can perform delivery and sensing, but with different roles: couriers prioritize timely deliveries and passively collect sensing data along the way, whereas RVs prioritize sensing to maximize coverage in each period (e.g., hourly). When courier capacity is insufficient, RVs assist with deliveries to reduce overdue orders. We also incorporate weather and road-network conditions, which affect agent behavior, travel speed, and region accessibility.

The *Simulator* models agent operations from time slot t to $t+1$ on a grid-based map. As shown in Fig. 2, we capture three behavior patterns: (1) Couriers always take reward-maximizing routes (modeled as the fastest routes), reflecting their preference for minimum travel time, and sense

passively; (2) RVs without delivery tasks follow system-generated sensing actions to maximize coverage; and (3) RVs with delivery tasks follow sensing actions under deadlines, balancing coverage and timely delivery.

B. Efficient Two-layer Decision-making

In this section, we present the details of the proposed efficient two-layer decision-making strategy with human-robot collaboration, which comprises two coupled layers: (i) the **upper layer**, an efficient MapReduce-based K-submodular maximization module for order dispatch (KSubMR), and (ii) the **lower layer**, a Deep Submodular Reward Q-Network (DSRQN) for route planning. Both couriers and RVs accept order dispatch actions, but only RVs strictly follow system-generated route planning actions.

1) *Optimization Objectives*: At the beginning of each time slot t , there are n_t^c available human couriers and n_t^{rv} available RVs. Let $D_t = \bigcup_{i=1}^{n_t^c + n_t^{rv}} d_i$ denote the set of available agents and O_t denote the set of pending orders at time t .

In our two-layer human-robot collaboration framework, the upper layer uses the order dispatch policy F_o (implemented as KSubMR) to assign orders to agents by maximizing delivery and sensing rewards, while the lower layer uses the route-planning policy F_s (implemented as DSRQN) to maximize RV sensing rewards subject to on-time delivery. Moreover, F_o relies on the sensing value estimated by F_s . As a result, F_o outputs the dispatch action a_t^o (i.e., the matching set S_t) and F_s outputs the routing actions a_t^s for RVs, forming the joint action $a_t = (a_t^o, a_t^s)$ in our MDP.

The upper-layer function F_o is defined as a binary assignment indicator:

$$F_o : O_t \times D_t \rightarrow \{0, 1\}, \quad (3)$$

where $F_o(o, d) = 1$ indicates that order $o \in O_t$ is assigned to agent $d \in D_t$. The resulting assignment set is

$$S_t = \{(o, d) \in O_t \times D_t \mid F_o(o, d) = 1\}. \quad (4)$$

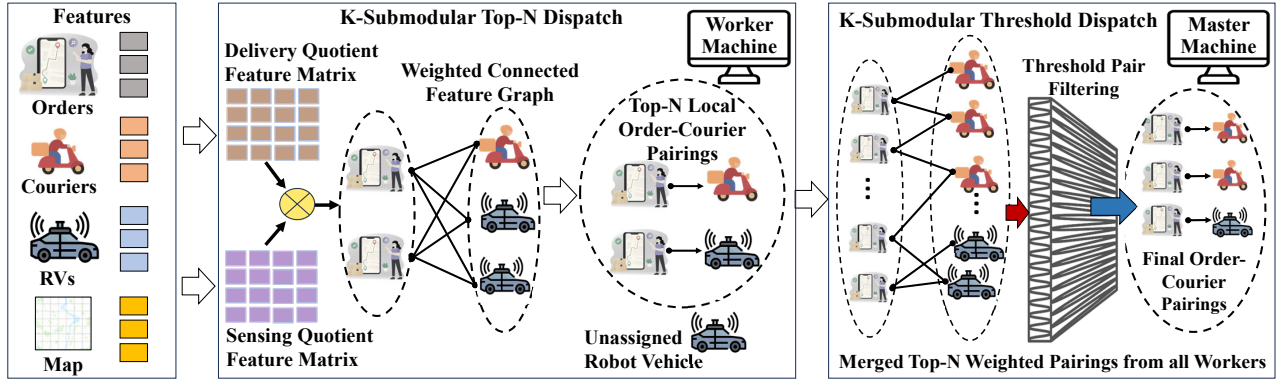


Fig. 3: An Overview of the Proposed KSubMR

we score each candidate pair (o, d) by

$$r(o, d) = r_d(o, d) + v_s(o, d), \quad (5)$$

where $o \in O_t$ and $d \in D_t$. Here, $r_d(o, d)$ is the delivery reward for completing order o (computed following Eq. 1), and $v_s(o, d)$ is the urban sensing value calculated in Section III-B.3. The resulting reward-driven many-to-many matching is formulated as the following optimization problem:

$$\begin{aligned} S_t &\leftarrow \arg \max_{S_t \subseteq O_t \times D_t} R_d(S_t) + V_s(S_t) \\ \text{s.t.} \quad &\sum_{d \in D_t} \mathbf{1}_{(o,d) \in S_t} \leq 1, \quad \forall o \in O_t, \\ &\sum_{o \in O_t} \mathbf{1}_{(o,d) \in S_t} \leq c_d, \quad \forall d \in D_t, \end{aligned} \quad (6)$$

where

$$\begin{aligned} R_d(S_t) &= \sum_{(o,d) \in S_t} r_d(o, d), \\ V_s(S_t) &= f_{\text{sub}} \left(\sum_{(o,d) \in S_t} v_s(o, d) \right). \end{aligned} \quad (7)$$

Here, $R_d(S_t)$ is the delivery reward of the assignment set S_t , $V_s(S_t)$ is its overall sensing value, and $f_{\text{sub}}(\cdot)$ is a submodular aggregation function. A key challenge in optimizing F_o in Eq. 3 is **asynchronous feedback**: $R_d(S_t)$ is available once S_t is decided, whereas $V_s(S_t)$ cannot be observed at dispatch time because it depends on subsequent RV routing actions. Thus, the objective in Eq. 6 is coupled with future decisions, making F_o nontrivial to optimize.

To address this issue, we propose UrbanHuRo with **hybrid reward-value feedback**. Specifically, the upper layer KSubMR applies the order dispatch policy F_o to solve Eq. 6 by combining the immediate reward $R_d(S_t)$ with the estimated value $V_s(S_t)$. The sensing value terms $v_s(o, d)$ and $V_s(S_t)$ are computed in the lower layer by DSRQN, and are fed back to KSubMR for dispatch optimization.

2) *KSubMR for Order Dispatch Decision-making*: The multi-objective optimization in Eq. 6 can be formulated as a weighted bipartite graph for many-to-many matching, for which the Kuhn–Munkres algorithm [6] is a common solution. However, as the numbers of orders, couriers, and RVs grow substantially, its *computational complexity* increases sharply. Coupled with the computational cost of deep RL, this approach becomes impractical for meeting the real-time requirements of order dispatch systems.

To address the computational-efficiency challenge and enable real-time coordination in highly dynamic environments, we propose the KSubMR algorithm, which integrates a two-round MapReduce framework with k -submodular maximization, as illustrated in Fig. 3. This approach leverages the submodular property [7] of the dispatch policy function F_o , which provides provable approximation guarantees in combinatorial optimization [8]. Moreover, the composability of local solutions makes it well-suited for parallelization in distributed frameworks such as MapReduce [9].

In the first MapReduce round, each *worker machine* runs the K-Submodular Top- N Dispatch algorithm (Algorithm 1) on its partition of orders and agent states $\{s_t^i\}$ to produce a local assignment set S_j . Here, M denotes the number of parallel workers. In the second MapReduce round, the *master machine* aggregates the worker outputs, reconstructs a bipartite graph over feasible order-agent pairs, and applies the K-Submodular Threshold Dispatch algorithm (Algorithm 2) to generate the final dispatch.

K-Submodular Top- N Dispatch (Algorithm 1): During the first MapReduce round, orders are distributed to worker machines with agent states s_t^i . Two quotient matrices are constructed, namely the Delivery Quotient (DQ) and Sensing Quotient (SQ), based on $r_d(o, d)$ and $v_s(o, d)$. Guided by these matrices, each order iteratively selects its top- N agents in a round-robin manner. Agents are ranked by the computable dispatch reward $R_d(\cdot)$ and the sensing value $V_s(\cdot)$ from DSRQN, where $\Delta_{o,d}(S_j)$ in Algorithm 1 denotes the marginal gain of adding (o, d) to the local assignment set S_j under $R_d(\cdot) + V_s(\cdot)$. The k -submodular structure reduces redundancy in local solutions and yields disjoint assignments. Each worker then transmits its results to the master node.

K-Submodular Threshold Dispatch (Algorithm 2): In the second MapReduce round, the master machine aggregates the worker outputs and reconstructs a bipartite graph connecting orders and agents to regenerate the quotient matrices. The final dispatch is obtained by applying a decreasing sequence of thresholds, starting from a selected order-agent pair. For each threshold τ , the algorithm scans unassigned order-agent pairs and selects those whose marginal gains exceed τ . This process continues until all orders are assigned or the threshold falls below a predefined value.

Algorithm 1 K-Submodular Top- N Dispatch

1: **Input:** dispatch reward function $R_d(\cdot)$, sensing value function $V_s(\cdot)$, agent limit N per order
2: **Output:** a sparse pair set $S \subseteq O \times D$
3: $S \leftarrow \emptyset$
4: **for** each $o \in O$ **do**
5: $S_j \leftarrow \emptyset$
6: **for** $i = 1$ to N **do**
7: $(o, d) \leftarrow \arg \max_{d \in D} \Delta_{o,d}(S_j)$
8: $S_j \leftarrow S_j \cup \{(o, d)\}$
9: **end for**
10: $S \leftarrow S \cup S_j$
11: **end for**
12: **return** S

Algorithm 2 K-Submodular Threshold Dispatch

1: **Input:** value oracle for monotone k -submodular $f_{\text{sub}}, O, D, \{c_d\}_{d \in D}, \epsilon \in (0, 1)$
2: **Output:** $S \subseteq O \times D$
3: $S \leftarrow \emptyset, \{S_d\}_{d \in D} \leftarrow \emptyset$
4: $\tau \leftarrow \tau_0 \leftarrow \max\{r_d(o, d) : o \in O, d \in D\}$
5: **while** $\tau > \epsilon \tau_0$ **do**
6: **for all** $(o, d) \in O \times D$ **do**
7: **if** $(o, d) \notin S$ **and** $f_{\text{sub}}(S \cup \{(o, d)\}) - f_{\text{sub}}(S) \geq \tau$ **and** $|S_d| < c_d$ **then**
8: $S_d \leftarrow S_d \cup \{(o, d)\}, S \leftarrow S \cup \{(o, d)\}$
9: **end if**
10: **end for**
11: $\tau \leftarrow (1 - \epsilon)\tau$
12: **end while**
13: **return** S

3) *DSRQN for Urban Sensing Decision-making:* We now present our Deep Submodular Reward Q-Network (DSRQN), which serves two objectives: (1) generating RV route-planning actions a_t^i to improve sensing coverage, and (2) estimating the aggregated sensing value $V_s(S_t)$ for the order dispatch layer. Specifically, DSRQN builds on a Q-value-based RL algorithm to estimate the expected return of routing actions and guide RV behavior.

Firstly, we describe how route-planning actions are generated. We introduce key concepts in DSRQN. The Q-value for a state-action pair, denoted as $Q(s_t^i, a_t^{s,i})$, represents the expected cumulative sensing reward when RV $r v^i$ takes routing action $a_t^{s,i}$ at state s_t^i . Accordingly, we define

$$Q(s_t^i, a_t^{s,i}) = \mathbb{E} \left[\sum_{k=0}^H r_{s,t+k}^i \mid s_t^i, a_t^{s,i} \right]. \quad (8)$$

In DSRQN, we parameterize the Q-function as $Q(s_t^i, a_t^{s,i}; \theta)$, where θ denotes the parameters of F_s . The sensing reward $r_{s,t}^i$ in Eq. 2 serves as the learning signal for estimating $Q(s_t^i, a_t^{s,i})$ and consists of three components: the regional reward $r_{g_t^i, d^i}^{\text{reg}}$, the neighboring reward $r_{g_t^i, d^i}^{\text{nbr}}$, and the timeout penalty $r_{g_t^i, d^i}^{\text{pen}}$.

The optimal Q-value $Q^*(s_t^i, a_t^{s,i})$ represents the maximum expected return achievable from state s_t^i by taking action $a_t^{s,i}$ and then making optimal routing decisions thereafter. In DSRQN, we learn a parameterized Q-function $Q(s_t^i, a_t^{s,i}; \theta)$ and derive the route-planning policy function F_s via a greedy action-selection rule:

$$F_s(s_t^i; \theta) = \arg \max_{a_t^{s,i}} Q(s_t^i, a_t^{s,i}; \theta). \quad (9)$$

Finally, F_s outputs the route-planning action $a_t^{s,i}$ for each RV $r v^i$ in the lower decision layer.

Next, we describe the computation of the aggregated sensing value $V_s(S_t)$, which is passed to KSubMR for solving F_o . We first define the sensing value $v_s(o, d)$ for an order-agent pair $(o, d) \in S_t$, which estimates the sensing gain of assigning order o to agent d (an RV in our setting, since only RVs follow system-generated routing actions). For illustration, we use an order o^j and an agent d^i as a running example. We adopt an RL value-function view to quantify the expected future sensing return from the agent's current state under the learned routing policy.

Since sensing is location-dependent, we define the region-level sensing value V_{reg} at region g given agent state s_t^i as

$$V_{\text{reg}}(g \mid s_t^i) = Q(s_t^i, F_s(s_t^i; \theta); \theta), \quad (10)$$

where F_s is the route-planning function in Eq. 9. Given an order $o_t^j \in O_t$ and an agent $d^i \in D_t$, let g_t^i be the current region of d^i and $o_{d,t}^j$ the drop-off region of o_t^j . We estimate the sensing value of assigning o_t^j to d^i as

$$v_s(o_t^j, d^i) = V_{\text{reg}}(o_{d,t}^j \mid s_t^i) - V_{\text{reg}}(g_t^i \mid s_t^i). \quad (11)$$

This value is computed as the difference between the region-level values at the drop-off region and the current region.

After obtaining $v_s(o_t^j, d^i)$ for each selected pair $(o_t^j, d^i) \in S_t$, we aggregate them to compute $V_s(S_t)$ in Eq. 7. To capture diminishing returns caused by spatial redundancy, we adopt a submodular aggregation function $f_{\text{sub}}(\cdot)$ and define

$$V_s(S_t) \triangleq f_{\text{sub}}(S_t) = \left(\sum_{(o_t^j, d^i) \in S_t} v_s(o_t^j, d^i) \right) \cdot (1 - \rho(S_t)), \quad (12)$$

where the redundancy penalty $\rho(S_t)$ is

$$\rho(S_t) = \frac{1}{|S_t|(|S_t| - 1)} \sum_{m \neq n} \exp \left(- \frac{\text{dis}((o^m, d^m), (o^n, d^n))}{2\sigma^2} \right). \quad (13)$$

Here, $\rho(\cdot)$ measures spatial redundancy among the paths induced by the assignments in S_t [10]. The term $\text{dis}((o^m, d^m), (o^n, d^n))$ denotes the distance between the m -th and n -th order-agent assignments, and the exponential kernel assigns higher similarity to closer (more redundant) paths. The normalization by $|S_t|(|S_t| - 1)$ keeps $\rho(S_t)$ scale-invariant, while the bandwidth σ controls the effective influence radius. The submodularity of $f_{\text{sub}}(\cdot)$ originates from the normalized pairwise similarity measure $\rho(\cdot)$, which induces diminishing returns in the aggregation in Eq. 12.

IV. EVALUATION

In this section, we conduct extensive experiments to evaluate the performance of our UrbanHuRo.

A. Evaluation Setup

1) *Parameter Setting and Experiment Environment:* We evaluate UrbanHuRo under fine-grained **spatial** and **temporal** settings to improve practicality. Specifically, we set the region size (GS) to $(0.05 \text{ km})^2$ and use a 1-minute time slot to support timely delivery and high-resolution sensing. We run experiments on 40 Intel(R) Xeon(R) Gold 5218R CPU @ 2.10GHz processors. We implement DSRQN in TensorFlow 2.4.1 and Python 3.8, and train on an NVIDIA TITAN RTX (TU102) GPU using real-world data distributions. Training uses an experience replay buffer of size 10^6 and a batch size of 128. We optimize Q-networks with Adam (learning rate 0.01) and set the discount factor γ to 0.9.

2) *Evaluation Metrics:* We evaluate the joint optimization of order dispatch and urban sensing using three metrics: number of overdue orders, sensing coverage, and the average courier income. These metrics are defined as follows:

- **Overdue Orders:** The number of orders not delivered by their deadlines during an hour.
- **Normalized Sensing Coverage (C_N):** The ratio of the number of visited locations achieved by a model to that achieved by UrbanHuRo.
- **Average Courier Income:** Average hourly income of human couriers, calculated based on total revenue from assigned orders minus overdue penalties.

3) *Datasets:* We utilize a real-world dataset from a food delivery platform for evaluation [11]. The dataset includes over 160K on-demand food delivery orders collected within a week in Shanghai, with around 2,200 active couriers at noon. It contains detailed information such as user IDs, pickup/drop-off locations with timestamps, delivery deadlines, and fees.

4) *Baselines:* We compare UrbanHuRo with the following five state-of-the-art baselines.

- **Fastest-Delivery (FastD):** This approach always selects the nearest agent and performs RV sensing using the Shortest Diversified Path Routing (SDPR) method, based on the K-Shortest Diversified Path Routing algorithm [12] with $k = 1$.
- **High-Sensing (HighS):** This approach selects the agent that provides the maximum coverage of unvisited locations along its route. It performs sensing using the Trajectory Scheduling for Maximum Task Coverage (TSMTC) algorithm [13], which selects the neighboring location with the highest coverage potential.
- **Joint-Delivery and Sensing Efficiency (JointDS)** [14]: This approach selects the agent with the highest ratio of sensed regions to driving distance. It utilizes the REASSIGN [15] routing method for RVs.
- **LSTAlloc** [16]: This approach aims to maximize courier profit by prioritizing order assignments based on a

balance between current income and driving distance. Any unassigned orders are then allocated to RVs.

- **AJRP** [17]: This approach estimates a cost-to-go function based solely on sensing gain, considering human couriers first. If no couriers are available, it applies the same strategy to assign active orders to RVs.

B. Evaluation Results

TABLE I: Overall Comparison.

# RVs	500	1000	2000	4000	500	1000	2000	4000	500	1000	2000	4000
Models	Overdue Orders				Coverage (C_N)				Courier Income			
HighS	291	287	284	281	1.01	0.79	0.72	0.81	5.76	5.55	5.07	5.00
FastD	16	15	14	13	0.05	0.03	0.02	0.02	5.63	4.39	3.14	2.64
JointDS	386	374	361	356	0.94	0.76	0.52	0.50	4.48	4.03	3.64	3.17
LSTAlloc	266	266	266	265	0.83	0.59	0.46	0.40	8.15	8.15	8.15	8.15
AJRP	390	374	366	361	0.88	0.62	0.48	0.42	5.71	5.69	5.61	5.48
UrbanHuRo	31	30	28	28	1.00	1.00	1.00	1.00	11.34	11.35	11.34	11.36

We provide an overall comparison of our experimental results in Table I. As shown, UrbanHuRo consistently achieves strong performance across different RV fleet-size settings in all the three metrics. The details are shown as below.

1) *Overdue Orders:* Fig. 4 compares the number of overdue orders at different times of the day, demonstrating the effectiveness of UrbanHuRo in minimizing delayed deliveries under dynamic, high-demand conditions. UrbanHuRo consistently matches FastD in maintaining timely fulfillment, even during peak periods. In contrast, other baselines exhibit noticeable spikes in overdue orders, particularly around noon, suggesting limited scalability for real-time dispatch.

Quantitatively, UrbanHuRo achieves an orders-of-magnitude reduction in overdue orders compared to state-of-the-art baselines, especially during core business hours (8:00 AM–8:00 PM) when system load is highest. This indicates that UrbanHuRo prioritizes time-sensitive dispatch without sacrificing sensing coverage. Unlike HighS and JointDS, which favor coverage at the expense of on-time delivery, UrbanHuRo strikes a better balance, making it more suitable for practical joint delivery and sensing.

2) *Sensing Coverage:* As shown in Table I, with 500 RVs, HighS shows similar map coverage to UrbanHuRo. However, as the number of RVs increases, UrbanHuRo outperforms all baselines. Notably, with 1000, 2000, and 4000 RVs, UrbanHuRo exceeds the best-performing baseline, HighS, by an average of 29.7% in coverage. This demonstrates UrbanHuRo’s scalability and adaptability in effectively balancing urban sensing with delivery efficiency.

3) *Courier Income:* Figure 5 shows that UrbanHuRo consistently boosts earnings for human couriers. On average, UrbanHuRo more than doubles courier income compared to HighS and JointDS, and exceeds LSTAlloc, a SOTA method specifically designed to maximize courier earnings, by about 39.2% in average. This improvement stems from UrbanHuRo’s ability to reduce overdue orders, which would otherwise diminish courier compensation. Importantly, UrbanHuRo achieves these income gains while maintaining low overdue orders and strong sensing coverage.

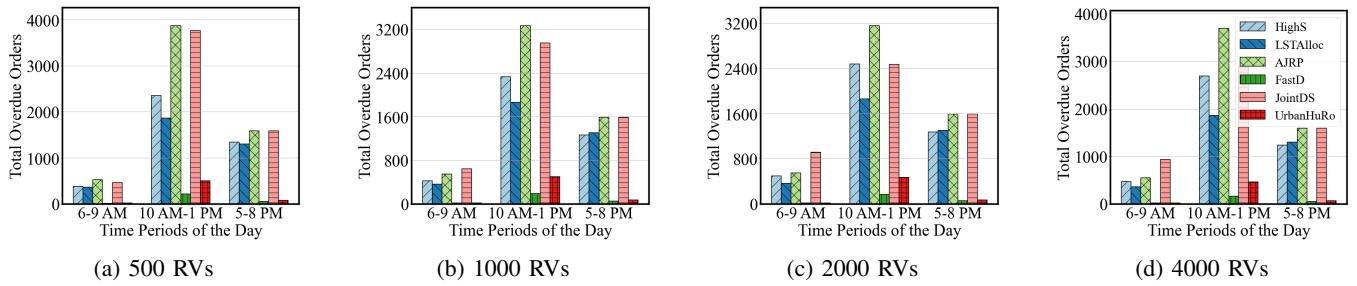


Fig. 4: Overdue orders across three key time periods, i.e., morning rush hours, noon peak hours, and evening rush hours.

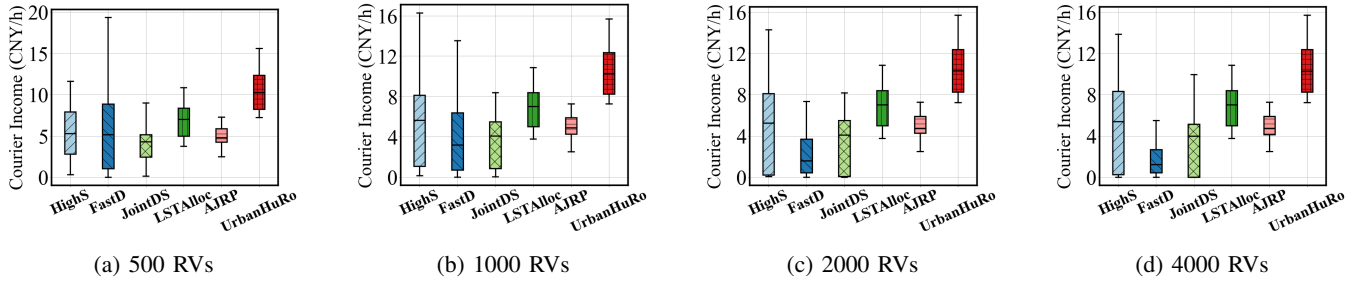


Fig. 5: Average hourly income of couriers under different numbers of RVs.

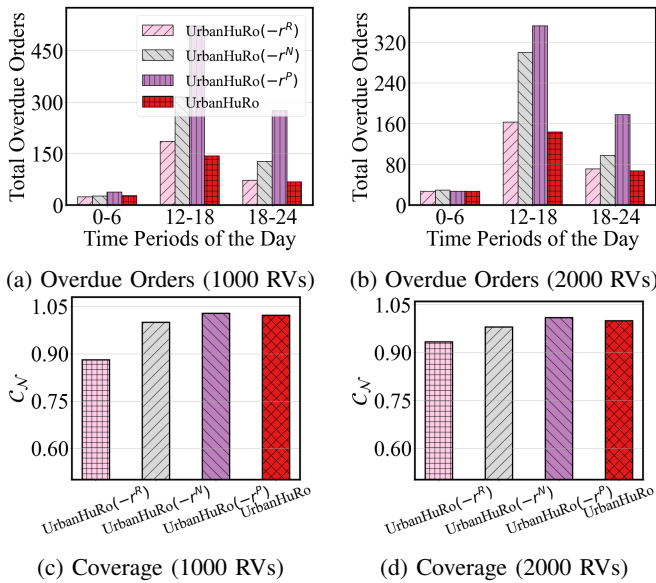


Fig. 6: Ablation Studies

4) *Ablation Study*: We evaluate three variants of UrbanHuRo: (i) UrbanHuRo(- r^R), excluding the regional reward $r_{g_t^i, d_t^i}^{\text{reg}}$; (ii) UrbanHuRo(- r^N), excluding the neighboring reward $r_{g_t^i, d_t^i}^{\text{nbr}}$; and (iii) UrbanHuRo(- r^P), excluding the timeout penalty $r_{g_t^i, d_t^i}^{\text{pen}}$. In Figures 6a and 6b, we observe that removing any of the three components leads to an increase in overdue orders. Notably, omitting the timeout penalty $r_{g_t^i, d_t^i}^{\text{pen}}$ causes the most significant rise in overdue orders throughout the day, whereas removing the neighboring reward $r_{g_t^i, d_t^i}^{\text{nbr}}$ has minimal impact. In Figures 6c and 6d, we also examine the impact of the three components on sensing coverage. The results show that removing either the regional reward $r_{g_t^i, d_t^i}^{\text{reg}}$ or the neighboring reward $r_{g_t^i, d_t^i}^{\text{nbr}}$ weakens the

system's sensing capability. In contrast, omitting the timeout penalty $r_{g_t^i, d_t^i}^{\text{pen}}$ enhances coverage, as the system no longer needs to sacrifice sensing gains to meet delivery deadlines.

V. RELATED WORK

A. Urban Service Optimization

Urban service optimization has attracted much interest from both the research community and industry due to its importance in reducing operational costs, increasing profits, and improving the quality of life. Recent studies have investigated effective order dispatch in ridesharing and logistics [3], [11], [18]–[20], route planning for urban sensing [21], [22], next location recommendation [23], and resource allocation and management in mobility systems [24]–[28]. While these works achieve substantial gains within individual services, they often overlook the potential synergies and reciprocal benefits that could be realized through the joint optimization of heterogeneous urban services.

B. Human-Robot Collaboration in Smart Cities

With the rapid advancement of smart cities, human-robot collaboration has become a key research focus for diverse urban services. For instance, in smart mobility and transportation, prior studies have examined interactions between delivery robots and human drivers in public transit systems [29]. In urban sensing and inspection, robots have been utilized to collaborate with municipal workers and local residents to collect sensing data [30]. Researchers have explored human-robot collaboration for public safety and emergency response [31]. Furthermore, human-robot collaboration has also been applied in social navigation within public spaces [32]. However, little research has explored human-robot collaboration for joint optimization of heterogeneous urban services, which is the focus of our work.

VI. CONCLUSION

This paper presents UrbanHuRo, a two-layer human-robot collaboration framework designed for joint optimization of heterogeneous urban services (e.g., crowdsourced delivery and urban sensing). UrbanHuRo comprises two key components: (i) a scalable distributed MapReduce-based K-Submodular maximization module called KSubMR for real-time order dispatch, and (ii) a deep submodular reward Q-network algorithm named DSRQN for dynamic urban sensing. These two modules are *interconnected* as the DSRQN will also provide sensing values for the KSubMR to optimize and balance real-time delivery and sensing performance. Our evaluation on real-world datasets from a food delivery platform demonstrates that UrbanHuRo improves sensing coverage by an average of 29.7% (with 1,000 to 4,000 RVs) and increases courier income by 39.2% on average, while substantially reducing overdue orders.

ACKNOWLEDGMENT

We sincerely thank all anonymous reviewers for their valuable suggestions. This work is partially supported by Florida State University, National Science Foundation under Grant Numbers 2411152, 2441179, 2231523, and 2140346.

REFERENCES

- [1] "Uber eats," <https://www.ubereats.com/>, Uber, 2026.
- [2] "Google street view," <https://www.google.com/streetview/>, Google, 2026.
- [3] L. Jiang, Y. Yang, and G. Wang, "Hcride: harmonizing passenger fairness and driver preference for human-centered ride-hailing," in *Proceedings of the Thirty-Fourth International Joint Conference on Artificial Intelligence*, 2025, pp. 10289–10297.
- [4] G. Wang, S. He, L. Jiang, S. Wang, F. Miao, F. Zhang, Z. Dong, and D. Zhang, "Fairmove: A data-driven vehicle displacement system for jointly optimizing profit efficiency and fairness of electric for-hire vehicles," *IEEE Transactions on Mobile Computing*, vol. 23, no. 6, pp. 6785–6802, 2023.
- [5] E. Wang, W. Liu, W. Liu, Y. Yang, B. Yang, and J. Wu, "Spatiotemporal urban inference and prediction in sparse mobile crowdsensing: A graph neural network approach," *IEEE Transactions on Mobile Computing*, 2022.
- [6] B. Guo, S. Wang, Y. Ding, G. Wang, S. He, D. Zhang, and T. He, "Concurrent order dispatch for instant delivery with time-constrained actor-critic reinforcement learning," in *2021 IEEE Real-Time Systems Symposium (RTSS)*. IEEE, 2021, pp. 176–187.
- [7] G. L. Nemhauser, L. A. Wolsey, and M. L. Fisher, "An analysis of approximations for maximizing submodular set functions—i," *Mathematical programming*, vol. 14, pp. 265–294, 1978.
- [8] B. Mirzasoleiman, A. Karbasi, R. Sarkar, and A. Krause, "Distributed submodular maximization," *The Journal of Machine Learning Research*, vol. 17, no. 1, pp. 8330–8373, 2016.
- [9] J. Dean and S. Ghemawat, "Mapreduce: simplified data processing on large clusters," *Communications of the ACM*, vol. 51, no. 1, pp. 107–113, 2008.
- [10] H. Lin and J. Bilmes, "A class of submodular functions for document summarization," in *Proceedings of the 49th annual meeting of the association for computational linguistics: human language technologies*, 2011, pp. 510–520.
- [11] L. Jiang, S. Wang, B. Guo, H. Wang, D. Zhang, and G. Wang, "Faircod: A fairness-aware concurrent dispatch system for large-scale instant delivery services," in *Proceedings of the 29th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, 2023, pp. 4229–4238.
- [12] Y. Lai, Y. Xu, D. Mai, Y. Fan, and F. Yang, "Optimized large-scale road sensing through crowdsourced vehicles," *IEEE Transactions on Intelligent Transportation Systems*, vol. 23, no. 4, pp. 3878–3889, 2022.
- [13] G. Fan, H. Jin, Q. Liu, W. Qin, X. Gan, H. Long, L. Fu, and X. Wang, "Joint scheduling and incentive mechanism for spatio-temporal vehicular crowd sensing," *IEEE Transactions on Mobile Computing*, vol. 20, 2019.
- [14] C. Xiang, Y. Zhou, H. Dai, Y. Qu, S. He, C. Chen, and P. Yang, "Reusing delivery drones for urban crowdsensing," *IEEE Transactions on Mobile Computing*, vol. 22, no. 5, pp. 2972–2988, 2021.
- [15] N. S. Lesmana, X. Zhang, and X. Bei, "Balancing efficiency and fairness in on-demand ridesourcing," *Advances in neural information processing systems*, vol. 32, 2019.
- [16] C. Xiang, W. Cheng, C. Lin, X. Zhang, D. Liu, X. Zheng, and Z. Li, "Lstalloc: A driver-oriented incentive mechanism for mobility-on-demand vehicular crowdsensing market," *IEEE Transactions on Mobile Computing*, vol. 23, no. 4, pp. 3106–3122, 2023.
- [17] S. Liu and Z. Luo, "On-demand delivery from stores: Dynamic dispatching and routing with random demand," *Manufacturing & Service Operations Management*, vol. 25, no. 2, pp. 595–612, 2023.
- [18] H. Zhang, G. Wang, X. Wang, Z. Zhou, C. Zhang, Z. Dong, and Y. Wang, "Nondbrem: Nondeterministic offline reinforcement learning for large-scale order dispatching," in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 38, no. 1, 2024, pp. 401–409.
- [19] S. Choudhury, K. Solovey, M. J. Kochenderfer, and M. Pavone, "Efficient large-scale multi-drone delivery using transit networks," in *2020 IEEE International Conference on Robotics and Automation (ICRA)*, 2020, pp. 4543–4550.
- [20] H. Makino and S. Ito, "Online multi-agent pickup and delivery with task deadlines," in *2024 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2024, pp. 8428–8434.
- [21] J. Rückin, F. Magistri, C. Stachniss, and M. Popović, "An informative path planning framework for active learning in uav-based semantic mapping," *IEEE Transactions on Robotics*, vol. 39, no. 6, pp. 4279–4296, 2023.
- [22] W. Ji, K. Han, and Q. Ge, "Route planning of mobile sensing fleets for repeatable environmental monitoring tasks," *Computers, Environment and Urban Systems*, vol. 119, p. 102285, 2025.
- [23] R. Xu, K. Cai, L. Jiang, Z. Hong, Y. Tian, and G. Wang, "Geogen: A two-stage coarse-to-fine framework for fine-grained synthetic location-based social network trajectory generation," in *The 40th Annual AAAI Conference on Artificial Intelligence (AAAI)*, 2026.
- [24] G. Wang, Y. Chen, S. Wang, F. Zhang, and D. Zhang, "Foretaxi: Data-driven fleet-oriented charging resource allocation in large-scale electric taxi networks," *ACM transactions on sensor networks*, vol. 19, no. 3, pp. 1–25, 2023.
- [25] L. Jiang, D. Yu, R. Xu, T. Tang, and G. Wang, "Uncertainty-aware predict-then-optimize framework for equitable post-disaster power restoration," in *Proceedings of the Thirty-Fourth International Joint Conference on Artificial Intelligence*, 2025, pp. 9719–9727.
- [26] H. Tan, Y. Yuan, H. Yan, S. Zhong, and Y. Yang, "Human preference-aware rebalancing and charging for shared electric micromobility vehicles," in *2024 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2024, pp. 9608–9615.
- [27] S. He, L. Pepin, G. Wang, D. Zhang, and F. Miao, "Data-driven distributionally robust electric vehicle balancing for mobility-on-demand systems under demand and supply uncertainties," in *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2020, pp. 2165–2172.
- [28] S. Ruan, J. Bao, Y. Liang, R. Li, T. He, C. Meng, Y. Li, Y. Wu, and Y. Zheng, "Dynamic public resource allocation based on human mobility prediction," *Proceedings of the ACM on interactive, mobile, wearable and ubiquitous technologies*, vol. 4, no. 1, pp. 1–22, 2020.
- [29] A. De Maio, G. Ghiani, D. Laganà, and E. Manni, "Sustainable last-mile distribution with autonomous delivery robots and public transportation," *Transportation Research Part C: Emerging Technologies*, vol. 163, p. 104615, 2024.
- [30] I. Munasinghe, A. Perera, and R. C. Deo, "A comprehensive review of uav-ugv collaboration: Advancements and challenges," *Journal of Sensor and Actuator Networks*, vol. 13, no. 6, p. 81, 2024.
- [31] T. H. Chung, V. Orekhov, and A. Maio, "Into the robotic depths: Analysis and insights from the darpa subterranean challenge," *Annual Review of Control, Robotics, and Autonomous Systems*, vol. 6, no. 1, pp. 477–502, 2023.
- [32] P. T. Singamaneni, P. Bachiller-Burgos, L. J. Manso, A. Garrell, A. Sanfeliu, A. Spalanzani, and R. Alami, "A survey on socially aware robot navigation: Taxonomy and future challenges," *The International Journal of Robotics Research*, vol. 43, no. 10, pp. 1533–1572, 2024.