

Review

Reinforcement learning for ridesharing: An extended survey[☆]Zhiwei (Tony) Qin^{a,*}, Hongtu Zhu^b, Jieping Ye^c^a Lyft Rideshare Labs, San Francisco, CA 94107, United States of America^b University of North Carolina, Chapel Hill, NC 27514, United States of America^c University of Michigan, Ann Arbor, MI 48109, United States of America

ARTICLE INFO

Keywords:

Ridesharing
Ride-pooling
Reinforcement learning
Multi-agent systems
Decision intelligence

ABSTRACT

In this paper, we present a comprehensive, in-depth survey of the literature on reinforcement learning approaches to decision optimization problems in a typical ridesharing system. Papers on the topics of rideshare matching, vehicle repositioning, ride-pooling, routing, and dynamic pricing are covered. Most of the literature has appeared in the last few years, and several core challenges are to continue to be tackled: model complexity, agent coordination, and joint optimization of multiple levers. Hence, we also introduce popular data sets and open simulation environments to facilitate further research and development. Subsequently, we discuss a number of challenges and opportunities for reinforcement learning research on this important domain.

1. Introduction

The emergence of ridesharing,² led by companies such as DiDi, Uber, and Lyft, has revolutionized the form of personal mobility. It is projected that the global rideshare industry will grow to a total market value of \$218 billion by 2025 (MarketsAndMarkets, 2018). However, how to improve operational efficiency is a major challenge for rideshare platforms, e.g., long passenger waiting time (Smith, 2019) and as high as 41% vacant time for ridesharing vehicles in a large city (Brown, 2020). The success of ridesharing, from the perspectives of the platforms, drivers, and passengers, requires sophisticated optimization of all the integrated components that collectively deliver the services.

Reinforcement learning (RL) is a machine learning paradigm that trains an agent to take optimal actions (measured by total cumulative reward) through interaction with the environment and getting feedback signals. It is a class of optimization methods for solving sequential decision-making problems with a long-term objective in a stochastic environment. Thanks to the rapid advancement in deep learning research and computing power, the integration of deep neural networks and RL has generated explosive progress in solving complex large-scale decision problems (Silver and Hassabis, 2016; Berner et al., 2019), attracting huge amount of renewed interests in the recent years. We are witnessing a similar trend in the ridesharing domain, where the demand and supply are highly stochastic and non-stationary, and the operational decisions are often sequential in nature and have strong spatiotemporal dependency. Simple greedy heuristics that only optimizes for immediate returns tend to produce short-sighted reactive policies, which do not align well with cumulative nature of the true performance measure of interest. The multi-step sequential nature of the decision-making (e.g., pricing, matching, repositioning) and the supply–demand stochasticity

[☆] This survey is a significantly expanded and refined version of Qin et al., (2021).

* Corresponding author.

E-mail address: zq2107@caa.columbia.edu (Z. Qin).

¹ The original version of this paper was completed while the author was at DiDi Labs, Mountain View, CA, USA. Any views presented in this paper do not represent those of either DiDi or Lyft.

² The business model covered in this survey is also often referred to as 'ride-sourcing' or 'on-demand ride services' (Wang and Yang, 2019).

in the environment pose enormous challenges to traditional predictive and optimization methods, spanning such issues as forecast accuracy, decision-time computational complexity, and adaptability to real-time changes. RL, on the other hand, presents itself as an excellent promising approach to these ridesharing optimization problems. RL methods are often highly data-driven, making them more suitable to situations where it is hard to build accurate predictive models. They are forward-looking, and yet, they do not explicitly depend on forecasting. And, by design, RL-based policies are dynamic and often light in decision-time complexity.

There are excellent surveys on RL for intelligent transportation (Haydari and Yilmaz, 2020; Yau et al., 2017), with in-depth coverage of traffic signals control and autonomous driving. Wang and Yang (2019) offers a broad review of ridesharing systems, whereas (Tong et al., 2020) surveys spatial crowdsourcing, which is a more general field than ridesharing. There has been no comprehensive review of the literature on RL for ridesharing, even though the field has attracted much attention and interest from the research communities for both RL and transportation just within the last few years (e.g., Shah et al., 2020; Tang et al., 2019; Al-Kanj et al., 2020; Shou and Di, 2020b). This paper aims to fill that gap by surveying the literature of this domain published in top conferences and journals in transportation (e.g., Transportation Research series, IEEE Transactions on Intelligent Transportation Systems, Transportation Science), data mining (e.g., KDD, ICDM, WWW, CIKM), and machine learning/AI (e.g., NeurIPS, AAAI, IJCAI). We describe the research problems associated with the various aspects of a ridesharing system, review the existing RL approaches proposed to tackle each of them, and discuss the challenges and opportunities. Reinforcement learning also has close relationship with approximate dynamic programming. Although it is not our goal in this paper to have a comprehensive review of this class of methods for problems in ridesharing, we aim to point the readers to the representative works so that they can refer to the further literature therein.

This survey is organized as follows: We lay out the ridesharing system architecture in Section 2 and define the scope of the problems to be reviewed. Within this section, as well as the subsequent sections of the survey, we clarify and draw connections among the different names that the problems are referred to, which are often due to the different communities that the authors are from and are easy to confuse by researchers new to ridesharing. In Section 3, we provide a concise review of the RL basics and the major algorithms adopted by the works in this survey. We review in details in Section 4 the literature for each problem described in Section 2 and the relevant data sets and environments. Finally in Section 5, we discuss some challenges and opportunities that we feel crucial in advancing RL for ridesharing.

2. Ridesharing

We first describe the architecture of a ridesharing system in this section, followed by explanation and clarification on the scopes of the problem associated with each module.

2.1. Architecture

A ridesharing service, in contrast to taxi hailing, matches passengers with drivers of vehicles for hire using mobile apps. In a typical mobile ridesharing system, there are five major modules: pricing, matching, repositioning, pooling, and routing. Fig. 1 illustrates the process and decision modules. When a potential passenger submits a trip request, the *pricing* module offers a quote, which the passenger either accepts or rejects. Upon acceptance, the *matching* module attempts to assign the request to an available driver. Depending on driver pool availability, the request may have to wait in the system until a successful match. Pre-match cancellation may happen during this time. The assigned driver then travels to pick up the passenger, during which time post-match cancellation may also occur. The pick-up location is usually where the passenger is making the request or he/she specifies. In some cases, it could be a public designated area, e.g., outside an airport or train station. After the driver successfully transports the passenger to the destination, she receives the trip fare and becomes available again. The *repositioning* module guides idle vehicles to specific locations in anticipation of fulfilling more requests in the future. Following the reposition recommendations is usually on a voluntary basis unless it is an autonomous ridesharing setting. Hence, it is common that the platform offers incentives to drivers for completing the repositions. When each driver takes only one passenger request at a time, i.e. only one passenger shares the ride with the driver, this mode is more commonly called ‘ride-hailing’. Ridesharing can refer to both ride-hailing and ride-pooling.³ In the *ride-pooling* mode, multiple passengers with different trip requests can share one single vehicle, so the pricing, matching, repositioning, and routing problems are different from those for ride-hailing and require specific treatment, in particular, considering the passengers already on board. The *routing* module provides turn-by-turn guidance on the road network to drivers/vehicles either in service of a passenger request or performing a reposition. The goal is to guide the vehicle to its destination efficiently and safely.

2.2. Problem scopes

First, we start from the pricing module. Since the trip fare is both the price that the passenger has to pay for the trip and the major factor for the income of the driver, pricing decisions influence both demand and supply distributions through price sensitivities of users, e.g., the use of surge pricing during peak hours. This is illustrated by the solid arrows pointing from the pricing module to orders and idle vehicles respectively in Fig. 1. The pricing problem in the ridesharing literature is in most cases dynamic pricing, which adjusts trip prices in real-time in view of the changing demand and supply. The pricing modules sits at the upstream position

³ In this survey, we do not cover topics on hitch, in which the driver is on his/her own trip with a specific destination.

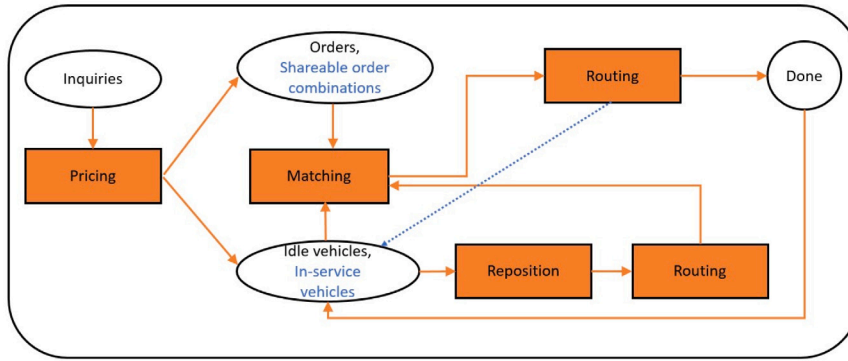


Fig. 1. The process flow of ridesharing operations. The solid orange rectangular boxes represent the modules described in Section 2, and the literature on the optimization problems associated with the modules are reviewed in the paper. The blue text and arrow apply exclusively to ride-pooling to account for the fact that order combinations and in-service vehicles are also eligible to participate in matching. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

with respect to the other modules and is a macro-level lever to achieve supply–demand (SD) balance. Although technically, driver pay can be determined by a separate module from pricing and has its own implication on supply elasticity and driver behavior, this paper follows the common setting where driver pay is closely associated (approximately proportional) to the trip fare so that pricing has the dual effect on demand and supply.

The ridesharing matching problem (Yan et al., 2020; Özkan and Ward, 2020; Qin et al., 2020) may appear under different names in the literature, e.g., order dispatching (Qin et al., 2020), trip-vehicle assignment (Bei and Zhang, 2018), and on-demand taxi dispatching (Tong et al., 2020). It is an online bipartite matching problem where both supply and demand are dynamic, with the uncertainty coming from demand arrivals, travel times, and the entrance-exit behavior of the drivers. Matching can be done continuously in a streaming manner or at fixed review windows (i.e., batching). Sophisticated matching algorithms often leverage demand prediction in some form beyond the actual requests, e.g., the value function in RL. Online request matching is not entirely unique to ridesharing. Indeed, ridesharing matching falls into the family of more general dynamic matching problems for on-demand markets (Hu and Zhou, 2022). A distinctive feature of the ridesharing problem is its spatiotemporal nature. A driver's eligibility to match and serve a trip request depends in part on her spatial proximity to the request. Trip requests generally take different amount of time to finish, and they change the spatial states of the drivers, affecting the supply distribution for future matching. The drivers and passengers generally exhibit asymmetric exit behaviors in that drivers usually stay in the system for an extended period of time, whereas passenger requests are lost after a much shorter waiting period in general.

Single-vehicle repositioning may refer to as taxi routing or passenger seeking in the literature. Taxi routing slightly differs in the setting from repositioning a rideshare vehicle in that a taxi typically has to be at a visual distance from the potential passenger to take the request whereas the matching radius of a mobile rideshare request is considerably longer, sometimes more than a mile. System-level vehicle repositioning, also known as driver dispatching, vehicle rebalancing/reallocation, or fleet management, aims to rebalance the global SD distributions by proactively dispatching idle vehicles to different locations. Repositioning and matching are similar to each other in that both relocate a vehicle to a different place as a consequence. In theory, one can treat repositioning as matching a vehicle to a virtual trip request, the destination of which is that of the reposition action, so that both matching and repositioning can be solved in a single problem instance. Typically in practice, these two problems are solved separately because they are separate system modules on most ridesharing platforms with different review intervals and objective metrics among other details.

The routing module described in Section 2.1 performs route guidance, which could be *dynamic routing* or *route planning* depending on the decision points. Dynamic routing is also called *dynamic route choice*, and route planning is alternatively referred to as the traffic assignment problem (Shou and Di, 2020a). In some cases, the reposition policy directly provides link-level turn-by-turn guidance with the goal of maximizing the driver's income, thus covering the role of dynamic routing albeit with a different objective. Dynamic routing is generally different from the *vehicle routing problem* (VRP) (Dantzig and Ramser, 1959). In VRP, the set of destinations that the vehicle has to visit is known in advance, and hence, it is a static problem. It mainly concerns with the sequence in which the destinations should be visited, considering the estimated travel time from point to point. In contrast, dynamic routing is associated with a road network, and the decision to make is which outgoing road (link) to follow at each intersection (node). The decision is adaptive to the changing traffic condition on the road network in real time. In the context of ride-pooling, there is another emerging problem in which the dynamic routing decisions with passenger(s) on board have to align with the overall objective of ride-pooling.

Some literature refers to the mode of multiple passengers sharing a ride as 'ridesharing'. In this paper, we use term '*ride-pooling*' (or 'carpool') to disambiguate the concept, as 'ridesharing' can refer to both single- and multiple-passenger rides. The seminal paper of Alonso-Mora et al. (2017a) shows that through ride-pooling optimization, the number of taxis required to meet the same trip demand can be significantly reduced with limited impact on passenger waiting times. In a pooling-enabled rideshare system, the matching, repositioning, and pricing modules all have to adapt to additional complexity. Compared to the regular ride-hailing

problem, the one with ride-pooling has considerably more complexity due to the more dynamic nature of the problem and the additional constraints and multiple objectives that have to be considered. In this case, the set of available vehicles are augmented, including both idle vehicles and occupied ones not at capacity. It is non-trivial to determine the set of feasible actions (one or more passengers to be assigned to a vehicle) for matching. Every time a new passenger is to be added to a non-idle vehicle, the route has to be recalculated using a VRP solver to account for the additional pick-up and drop-off, the travel times for all the passengers already on board are updated, and the vehicle capacity, the waiting time and detour distance constraints are checked. In-service routing in the context of ride-pooling is discussed in Section 4.5.

3. Reinforcement learning

We briefly review the RL basics and the major algorithms, especially those used in the works reviewed by this survey. For a complete reference, see, e.g., Sutton and Barto (2018).

3.1. Basics

RL is based on the Markov decision process (MDP) framework, where the agent (the decision-making entity) has a *state* s (e.g., the location of a vehicle) in the state space S and can perform an *action* a (e.g., to dispatch or idle) in the action space \mathcal{A} . The action is determined by the agent's policy, $\pi(s) : S \rightarrow \mathcal{A}$. If the policy is stochastic, then $\pi(a|s)$ gives the probability of selecting a given s . After executing the action, the agent receives an immediate *reward* $R(s, a)$ from the environment, and its state changes according to the transition probabilities $P(\cdot|s, a)$. The process repeats until a terminal state or the end of the horizon is reached, giving a sequence of the triplets $(s_t, a_t, r_t)_{t=0}^{t=T}$, where t is the epoch index, T is the final epoch at the terminal state or the end of the horizon, and r_t is a sample of R . The objective of the MDP is to maximize the cumulative reward over the horizon. A key quantity to compute is the value function associated with π ,

$$V^\pi(s) := E_\pi \left[\sum_{t=0}^{t=T} \gamma^t r_t \mid s_0 = s \right],$$

which satisfies the Bellman equation,

$$V^\pi(s_t) = \sum_{a_t} \pi(a_t|s_t) \sum_{s_{t+1}, r_t} P(s_{t+1}, r_t|s_t, a_t) \left(r_t(s_t, a_t) + \gamma V^\pi(s_{t+1}) \right). \quad (1)$$

Similarly, we have the action-value function,

$$Q^\pi(s, a) := E_\pi \left[\sum_{t=0}^{t=T} \gamma^t r_t \mid s_0 = s, a_0 = a \right],$$

which conditions on both s and a . The optimal state- and action-values are denoted by V^* and Q^* , evaluated at the optimal policy π^* . The objective of an MDP is to find an optimal policy π^* that maximizes the long-term cumulative discounted reward, i.e., $\pi^* := \arg \max_\pi E_\pi[V^\pi(s)]$.

3.2. Algorithms

Given P and R , which specifies the MDP, and π , we can compute V^π by iteratively applying the Bellman Eq. (1),

$$V^\pi(s) \leftarrow \sum_a \pi(a|s) \sum_{s', r} P(s', r|s, a) \left(r + \gamma V^\pi(s') \right). \quad (2)$$

This is called *policy evaluation*. Using policy evaluation as a sub-routine, we can again iteratively improve the policy through *policy iterations*, which generate a new policy at each outer iteration by acting greedily with respect to V^π ,

$$a^* \leftarrow \arg \max_a \sum_{s', r} P(s', r|s, a) \left(r + \gamma V^\pi(s') \right). \quad (3)$$

As a special instance, we can collapse the inner policy evaluation loop to a single iteration and compute $V^*(s), \forall s \in S$, by the *value iterations*,

$$V(s) \leftarrow \max_a \sum_{s', r} P(s', r|s, a) \left(r + \gamma V(s') \right). \quad (4)$$

If we estimate (P, R) from data, we have a basic model-based RL method. A model-free method learns the value function and optimizes the policy directly from data without learning and constructing a model of the environment. A common example is *temporal-difference (TD) learning* (Sutton, 1988), which iteratively updates V^π by TD-errors using π -generated trajectory samples and bootstrapping,

$$V^\pi(s) \leftarrow V^\pi(s) + \alpha \left(r + \gamma V^\pi(s') - V^\pi(s) \right), \quad (5)$$

where s' is the next state in the trajectory after s , α is the step size (or learning rate), and the term that it scales is the TD error. If learning the optimal action-values for control is the goal, we similarly have *Q-learning* (Watkins and Dayan, 1992), which updates the action-value function $Q(s, a)$ to approximate $Q^*(s, a)$ by

$$Q(s, a) \leftarrow Q(s, a) + \alpha \left(r + \gamma \max_{a'} Q(s', a') - Q(s, a) \right). \quad (6)$$

Q-learning is an *off-policy* algorithm, where the behavior policy, which collects the experience data and typically involves exploration, is different from the target policy that we are trying to learn and in this case, is the optimal policy. The *on-policy* counterpart of Q-learning is SARSA, which basically generalizes TD-learning (5) to the action-value function associated with the behavior policy π (same as the target policy):

$$Q^\pi(s, a) \leftarrow Q^\pi(s, a) + \alpha \left(r + \gamma Q^\pi(s', a') - Q^\pi(s, a) \right), \quad (7)$$

where a' is the action executed by the agent at state s' in the experience data. The *deep Q-network* (DQN) (Mnih et al., 2015) approximates $Q(s, a)$ by a neural network Q_w parametrized by w along with a few heuristic techniques like experience replay (Lin, 1992) and a target network to improve training stability. These techniques are critical to the successful of DQN in playing Atari games and many other applications, due to the deadly triad issue (Sutton and Barto, 2018) of reinforcement learning when one tries to combine bootstrapping (i.e., TD-learning and Q-learning), off-policy training (i.e., Q-learning), and function approximations (i.e., neural networks), which may lead to instability and divergence. The algorithms introduced so far are all value-based methods, which focus on learning the value function, and the policy is derived from the learned value function by, e.g., $\arg \max_a Q(s, a)$. Neural network-based value function approximation is important to ridesharing applications because the state is often high dimensional with the incorporation of SD contextual features. Tabular methods suffer from the curse of dimensionality and are not tractable in this case.

A policy-based method directly learns π (which is also called the *actor* and parametrized by θ) by performing stochastic gradient descent. The central step is computing the policy gradient (PG), the gradient of the cumulative reward $J(\theta)$ with respect to the policy parameters θ ,

$$\nabla_\theta J(\theta) = \sum_s \mu(s) \sum_a Q^\pi(s, a) \nabla_\theta \pi(a|s, \theta), \quad (8)$$

where μ is the on-policy distribution under π . A more common (equivalent) form of the PG (8) is

$$\sum_s \mu(s) \sum_a \pi(a|s, \theta) Q^\pi(s, a) \nabla_\theta \log \pi(a|s, \theta), \quad (9)$$

which most of the policy-based methods are based on. REINFORCE (Williams, 1992) is a classical PG method, which uses Monte Carlo (MC) rollout to obtains the sample-based update

$$\theta_{t+1} \leftarrow \theta_t + \alpha G_t \nabla_\theta \log \pi(a|s, \theta), \quad (10)$$

where G_t is an MC approximation of Q^π . As in a value-based method, we can also use function approximation for the action values. The function approximator (e.g., a neural network) Q_w is called the *critic*, and the resulting algorithm is an *actor-critic* (AC) method. It is well-recognized that the ‘baseline’ version of (8),

$$\sum_s \mu(s) \sum_a \pi(a|s, \theta) \left(Q^\pi(s, a) - b(s) \right) \nabla_\theta \log \pi(a|s, \theta), \quad (11)$$

where $b(s)$ is an action-independent baseline, reduces the variance in the sample gradient and helps speed up learning. Since a natural choice of such a baseline is the state value, the critic often learns the advantage function $Q(s, a) - V(s)$, and the method is called Advantage Actor Critic (A2C). The evaluation of an action is based on how good it can be with respect to the average over all actions, the benefit of which is to reduce the high variance in the actor and to stabilize the model. Mnih et al. (2016) extend A2C to an asynchronous version (A3C) where independent agents interact with their own copy of the environment and update their model parameters with the master copy asynchronously. This architecture enables much more efficient utilization of the CPU cores through parallel threads and hence accelerates the training. The proximal policy optimization (PPO) (Schulman et al., 2017) optimizes a clipping surrogate objective with respect to the advantage to promote conservative updates to π and is a popular choice of training algorithm for RL problems where policy-based methods are suitable (e.g., with continuous actions).

An MDP can be extended to a Markov game involving multiple agents to form the basis for multi-agent RL (MARL). Many MARL algorithms, e.g., Yang et al. (2018) and Lowe et al. (2017) focus on agent communication and coordination, in view of the intractable action space in MARL.

3.3. Approximate dynamic programming

A family of methods closely related to RL is approximate dynamic programming (ADP) (Powell, 2007) for solving stochastic dynamic programs (DP), of which the Bellman equation for MDP is an instance. In ADP methods, unlike that typically seen in RL, a post-decision state s_t^x is often defined to represent the intermediate state to which the current state s_t will transition deterministically given the action a_t before the random factors ω_t (e.g., demand appearance and cancellation) in the environment realize. With ω_t

fully realized, the state transitions into the next pre-decision state s_{t+1} . The value function in an ADP method is defined on the post-decision state and is approximated by a particular functional form. Given the approximated values, the original optimization problem is solved to obtain the decision solution for the current time step. Linear function approximation is popular (e.g., Simao et al., 2009; Yu and Shen, 2019; Al-Kanj et al., 2020) because the dual variables associated with the solution to the current-stage optimization can be used to update the linear function parameters. Then, the state is advanced to the next pre-decision state, and the iteration continues until convergence. By nature, ADP methods are on-policy methods. Recently, neural network-based value function approximation (Shah et al., 2020) has also been adopted and developed due to their higher level of flexibility. In this case, the value function updates largely follow the DQN scheme. The ADP methods for ridesharing reviewed in this survey solve system-level stochastic DP problems (e.g., matching and repositioning) and aim to approximate the system value by decomposing it into local or driver-centric values, and the update schemes employed fall into the family of approximate value iterations.

4. Reinforcement learning for ridesharing

We review the RL literature for ridesharing in this section grouped by the core operational problems described in Section 2. We first cover pricing, matching, repositioning, and routing in the context of ride-hailing. Then, we will review works on those problems specific to ride-pooling.

4.1. Pricing

RL-based approaches have been developed for dynamic pricing in one-sided retail markets (Raju et al., 2003; Bertsimas and Perakis, 2006), where pricing changes only the demand pattern per customers' price elasticity. The ridesharing marketplace, however, is more complex due to its two-sided nature and spatiotemporal dimensions. In this case, pricing is also a lever to change the supply (driver) distribution if price changes are broadcast to the drivers. Chen et al. (2021) describe examples of such elasticity functions for both demand and supply for their simulation environment.

The challenges in dynamic pricing for ridesharing lie in both its exogeneity and endogeneity. Dynamic pricing on trip inquiries changes the subsequent distribution of the submitted requests through passenger price elasticity. The requests distribution, in turn, influences future supply distribution as drivers fulfill those requests. On the other hand, the trip fares influence the demand for ridesharing services at given locations, and these changes will affect the pool of waiting passengers, which further affects the passengers' expected waiting times. Again, it will influence the demand either through cancellation of the current requests or the conversion of future trip inquiries. Because of its close ties to SD distributions, dynamic pricing is often jointly optimized with order matching or vehicle repositioning. Within the (non-RL) operations research literature, dynamic pricing for ridesharing has already been studied and analyzed in conjunction with matching (Yan et al., 2020; Özkan and Ward, 2020) and from the spatiotemporal perspective (Ma et al., 2020; Bimpikis et al., 2019; Hu et al., 2022), covering optimality and equilibrium analyses.

The complex interaction between pricing and the SD makes it hard to explicitly develop mathematical models that adapt well to dynamic and stochastic environments, and RL comes in as a promising direction to address these challenges by considering endogeneity and exogeneity as part of the environment dynamics.

Table 1 summarizes the reviewed works on RL for dynamic pricing in ridesharing. As one of the early RL works, Wu et al. (2016) consider a simplified ridesharing environment which captures only the two-sidedness of the market but not the spatiotemporal dimensions. The state of the MDP is the current price plus SD information. The action is to set a price, and the reward is the generated profit. A Q-learning agent is trained in a simple simulator, and empirical advantage in the total profit is demonstrated against other heuristic approaches. More recent works leverage the spatiotemporal nature of the pricing actions and take into account the spatiotemporal long-term values in the pricing decisions. Chen et al. (2019a) integrate contextual bandits and the spatiotemporal value network developed in Tang et al. (2019) for matching to jointly optimize pricing and matching decisions. In particular, the pricing actions are the discretized price percentage changes and are selected by a contextual bandits algorithm, where the long-term values learned by the value network are incorporated into the bandit rewards. In Turan et al. (2020), the RL agent determines both the price for each origin–destination (OD) pairs and the reposition/recharging decisions for each electric vehicle in the fleet. The state contains global information such as the electricity price in each zone, the passenger queue length for OD pair, and the number of vehicles in each zone and their energy levels. The reward accounts for trip revenue, penalty for the queues, and operational cost for charging and repositioning. Due to the multi-dimensional continuous action space, PPO is used to train the agent in a simulator. Song et al. (2020) perform a case study of ridesharing in Seoul. They use a tabular Q-learning agent to determine spatiotemporal pricing, and extensive simulations are performed to analyze the impact of surge pricing on alleviating the problem of marginalized zones (areas where it is consistently hard to get a taxi) and on improving spatial equity. Chen et al. (2021) adopt PPO to optimize the spatiotemporal pricing decisions for each hexagonal cell in terms of the per-km rate for the excess mileage beyond a base trip distance and the per-km rate for driver wage, for the objective of maximizing profits (revenue minus wage). The agent is modeled as a global decision-maker with state information of the numbers of open requests, vacant vehicles, occupied vehicles in each grid cell at time t and historical demand at time $t - 1$. Unlike the works above that focus on the pricing decisions, Mazumdar et al. (2017) study from a different perspective of the pricing problem. The proposed risk-sensitive inverse RL method (Ng et al., 2000) recovers the policies of different types of passengers (risk-averse, risk-neutral, and risk-seeking) in view of surge pricing. The policy determines whether the passenger should wait or take the current ride.

As discussed in Section 2.2, under the setting where the driver pay is associated with the trip fare, the dynamic pricing policy also affects supply elasticity, i.e., drivers' decisions on participation in a given marketplace, working hours, and in some cases, the

Table 1
Summary of literature for Pricing.

| Paper | Agent | State | Action | Reward | Algorithm | Environment |
|------------------------|---|--|--|--|--|--|
| Wu et al. (2016) | Global decision-maker | Current trip price (same for all trips), SD info | Price | Profit | Q-learning | No spatiotemporal dimensions |
| Chen et al. (2019a) | Global decision-maker | Features of the trip request | Discretized price change percentage | Profit | Contextual bandits with action values partly computed by CVNet | Ride-hailing simulator with pricing module and passenger elasticity model |
| Turan et al. (2020) | Global decision-maker for pricing and EV charging | Electricity price in each zone, passenger queue length for each OD pair, number of vehicles in each zone and their energy levels | Price for each OD pair, reposition/charging for each vehicle | Trip revenue - penalty for queues - operational cost for charging and reposition | PPO | Simulator |
| Song et al. (2020) | Global decision-maker | Location, time | Price for spatialtemporal grid cells | Trip price minus penalty for driver waiting | Q-learning | Case study: ride-hailing simulation of Seoul |
| Mazumdar et al. (2017) | Passenger | Price multiplier, time, if a ride has completed | Wait, take current ride | Trip price to pay | Risk-sensitive inverse RL | Historical data |
| (Chen et al., 2021) | Global decision-maker | Number of open requests, vacant vehicles, and occupied vehicles in each grid cell at time t , and demand in time $t - 1$ | Joint actions of price (per-km for excess mileage) and wage (per-km rate) for each grid cell | Profit: revenue minus wage | PPO | Simulation based on Hangzhou data from DiDi; modeling on both supply and demand elasticity |

probability of accepting a given assignment, depending on the rules of the particular ridesharing platform (Chen and Sheldon, 2016; Sun et al., 2019; Angrist et al., 2021). Although not yet being widely considered in RL approaches to pricing, supply elasticity is an important piece of system state information that has significant implication to the sequence of pricing decisions. For the closely related topic of driver incentives design, Shang et al. (2019, 2021) adopts a learning-based approach to construct a generative model for driver behavior with respect to the incentives policy and subsequently trains an RL agent to optimize the incentives design for system-level metrics. Perhaps this example sheds some light on how RL is able to help improve pricing policies in view of supply-side effects.

4.2. Online matching

The rideshare matching problem and its generalized forms have been investigated extensively in the field of operations research (see e.g., Özkan and Ward, 2020; Hu and Zhou, 2022; Lowalekar et al., 2018 and the references therein). Typically, both the open trip requests and available drivers are batched within time windows of fixed length as they arrive at the system, and they are matched at predefined discrete review times. See Fig. 2 for an illustration. Hence, ridesharing matching is an online stochastic problem (Qin et al., 2020). Outside the RL literature, Lowalekar et al. (2018) approach the problem through stochastic optimization and use Bender's decomposition to solve it efficiently. To account for the temporal dependency of the decisions, Hu and Zhou (2022) formulate the problem as a stochastic DP and propose heuristic policies to compute the optimal matching decisions. For a related problem, the truckload carriers assignment problem, Simao et al. (2009) also formulate a dynamic DP but with post-decision states so that they are able to solve the problem using ADP. In each iteration, a demand path is sampled, and the value function is approximated in a linear form and updated using the dual variables from the LP solution to the resulting optimization problem.

The RL literature for rideshare matching (see Table 2) typically aims to optimize the total driver income and the service quality over an extended period of time. Service quality can be quantified by *response rate* and *fulfillment rate*. Response rate is the ratio of the matched requests to all trip requests. Since the probability of pre-match cancellation is primarily a function of response time (pre-match waiting time), the total response time is an alternative metric to response rate. Fulfillment rate is the ratio of completed requests to all requests and is no higher than the response rate. The gap is due to post-match cancellation, usually because of the waiting for pick-up. Hence, the average pick-up distance is also a relevant quantity to observe. Fig. 3 shows the detailed flow of matching a single trip request together with the quantities discussed above.

In terms of the MDP formulation, driver agent is a convenient modeling choice for its straightforward definition of state, action, and reward, in contrast to system-level modeling where the action space is exponential. In this case, the rideshare platform is naturally a multi-agent system with a global objective. A common approach is to crowdsource all drivers' experience trajectories to




Fig. 2. The order matching process with batching from the system perspective (Qin et al., 2020). The assignments are for illustration only.

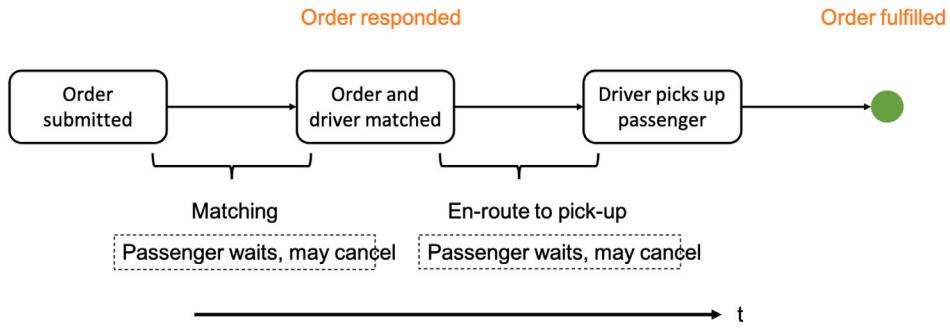


Fig. 3. The order matching process from a single request's perspective.

train a single agent and apply it to all the drivers to generate their matching policies (Xu et al., 2018; Wang et al., 2018; Tang et al., 2019). Since the system reward is the sum of the drivers' rewards, the system value function does decompose into the individual drivers' value functions computed by each driver's own trajectories. The approximation here is using a single value function learned from all drivers' data. See Qin et al. (2020) for detailed discussions. Specifically, Xu et al. (2018) learn a tabular driver value function using TD(0), and Wang et al. (2018), Tang et al. (2019), Holler et al. (2019) apply DQN-type of training to learn a value network. In particular, Tang et al. (2019) design a spatiotemporal state-value network using hierarchical coarse coding and cerebellar embedding memories for better state representation and training stability. Holler et al. (2019) develop an action-value network that leverages global SD information, which is embedded into a global context by attention.

This type of single-agent approach avoids dealing explicitly with the multi-agent aspect of the problem and the interaction among the agents during training. Besides simplicity, this strategy has the additional advantage of being able to easily handle a dynamic set of agents (and hence, a changing action space) (Ke et al., 2020b). On the other hand, order matching requires strong system-level coordination in that a feasible solution has to satisfy the one-to-one constraints. To address this issue, Xu et al. (2018) and Tang et al. (2019) use the learned state values to populate the edge weights of a bipartite assignment problem to generate a collective-greedy policy (Qin et al., 2020) with respect to the state values. Holler et al. (2019) assume a setting where drivers are matched or repositioned sequentially so that the policy output always satisfies the matching constraints.

Leveraging MARL, Li et al. (2019), Jin et al. (2019) and Zhou et al. (2019) directly optimize the multi-agent system. One significant challenge is scalability since any realistic ridesharing setting can easily involve thousands of agents, precluding the possibility of dealing with an exact joint action space. Li et al. (2019) apply mean-field MARL to make the interaction among agents tractable, by taking the 'average' action of the neighboring agents to approximate the joint actions. Zhou et al. (2019) argue that no explicit communication among agents is required for order matching due to the asynchronous nature of the transitions and propose independent Q-learning with centralized KL divergence (of the supply and demand distributions) regularization. Both Li et al. (2019) and Zhou et al. (2019) follow the centralized training decentralized execution paradigm. Jin et al. (2019) take a different approach treating each spatial grid cell as a worker agent and a region of a set of grid cells as a manager agent, and they adopt hierarchical RL to jointly optimize order matching and vehicle repositioning.

In practical settings, the online matching policy often has to balance among multiple objectives (Lyu et al., 2019), e.g., financial metrics and customer experience metrics. The rationale is that persistent negative customer experience will eventually impact long-term financial metrics as users churn the service or switch to competitors. There are two potential ways that one can leverage RL to approach this problem. The explicit approach is to directly learn a policy that dynamically adjusts the weights to combine the multiple objectives into a single reward function. With the abundance of historical experience data, inverse RL can be used to learn the relative importance of multiple objectives under a given unknown policy (Zhou et al., 2021). The implicit approach is to capture the necessary state signals that characterize the impact of the metrics not explicitly in the reward function, so that the learned value function correctly reflect the long-term effect of the multiple metrics. As discussed in Section 5.5, the long feedback loop is a potential challenge here.

Besides the driver-passenger pairing decisions, there are other important levers that can be optimized within the matching module, namely the matching window and the matching radius (Yang et al., 2020a). The matching window determines when to match a request (or a batch of requests). A larger window increases pre-match waiting time but may decrease pick-up time for matched requests because of more available drivers. There have been several RL works on the matching window optimization, which can be done from the perspective of a request itself (Ke et al., 2020b) or the system (Wang et al., 2019; Qin et al., 2021a). In Ke et al. (2020b), each trip request is an agent. An agent network is trained centrally using pooled experience from all agents to decide whether or not to delay the matching of a request to the next review window, and all the agents share the same policy. To encourage cooperation among the agents, a specially shaped reward function is used to account for both local and global reward feedback. They also modify the RL training framework to address the delayed reward issue by sampling complete trajectories at the end of training epochs to update the network parameters. Wang et al. (2019) take a system's view and propose a Restricted Q-learning algorithm to determine the length of the current review window (or batch size). They show theoretical analysis results on the performance guarantee in terms of competitive ratio for dynamic bipartite graph matching with adaptive windows. Qin et al. (2021a) take a similar modeling perspective but use the AC method with experience replay (ACER) (Wang et al., 2016) that combines on-policy updates (through a queuing-based simulator) with off-policy updates. The matching radius defines how far an idle driver can be from the origin of a given request to be considered in the matching. It can be defined in travel distance or time. A larger matching radius may increase the average pick-up distance but requests are more likely to be matched within a batch window, whereas a smaller radius renders less effective driver availability but it may decrease the average pick-up distance. Both the matching window and radius are trade-offs between pre-match and post-match waiting times (and hence, cancellation). So far, few effort through RL has been devoted to matching radius optimization. The joint optimization of the matching window and radius is certainly another interesting line of research.

Because of its generalizability, matching for ridesharing is closely related to a number of online matching problems in other domains, the RL methods to which are also relevant and can inform the research in rideshare matching. Some examples are training a truck agent using DQN with pooled experience to dispatch trucks for mining tasks (Zhang et al., 2020b), learning a decentralized value function using PPO with a shaped reward function for cooperation (in similar spirit as Ke et al. (2020b)) to dispatch couriers for pick-up services (Chen et al., 2019b), and designing a self-attention, pointer network-based policy network for a system agent to assign participants to tasks in mobile crowdsourcing (Shen et al., 2020).

4.3. Vehicle repositioning

Vehicle repositioning from a single-driver perspective (i.e., taxi routing) has a relatively long history of research since taxi service has been in existence long before the emergence of rideshare platforms. Likewise, research on RL-based approaches for this problem (see Table 3) also appeared earlier than that on system-level vehicle repositioning.

For the taxi routing problem, each driver is naturally an agent, and the objective thus focuses on optimizing individual reward. Common reward definitions include trip fare (Rong et al., 2016), net profit (income - operational cost) (Verma et al., 2017), idle cruising distance (Garg and Ranu, 2018), and ratio of trip mileage to idle cruising mileage (Gao et al., 2018). Earlier works (Han et al., 2016; Wen et al., 2017; Verma et al., 2017; Garg and Ranu, 2018) optimize the objective within a horizon up to the next successful match (i.e., A-to-B and A-to-C in Fig. 4), but it is now more common to consider a long-term horizon, where an episode usually consists of a trajectory over one day (Lin et al., 2018; Shou et al., 2020; Jiao et al., 2021). We illustrate these concepts in Fig. 4.

The type of actions of an agent depends on the physical abstraction adopted. A simpler and more common way of representing the spatial world is a grid system, square or hexagonal⁴ Han et al. (2016), Wen et al. (2017), Verma et al. (2017), Gao et al. (2018), Lin et al. (2018), Rong et al. (2016), Jiao et al. (2021) and Shou et al. (2020). In this setting, the action space is the set of neighboring cells (often including the current cell). Shou and Di (2020b) explain the justification for this configuration. Determination of the specific destination point is left to a separate process, e.g., pick-up points service (Jiao et al., 2021). The more realistic abstraction is a road network, in which the nodes can be intersections or road segments (Garg and Ranu, 2018; Yu et al., 2019; Zhou et al., 2018; Schmoll and Schubert, 2020). The action space is the adjacent nodes or edges of the current node. This approach supports a turn-by-turn guiding policy but requires more map information at run time.

Most of the papers adopt a tabular value function, so the state is necessarily low-dimensional, including spatiotemporal information and sometimes additional categorical statuses. Shou et al. (2020) have a boolean in the state to indicate if the driver is

⁴ The hexagonal grid system is the industry standard.

Table 2
Summary of literature for Online Matching.

| Paper | Agent | State | Action | Reward | Algorithm | Environment | Notes |
|--|--|--|---|--|--|---|--|
| Xu et al. (2018) | Driver | Assignment to a specific order or idle | Location, time | Trip price | Tabular TD(0) for learning state values offline + Hungarian method for generating the assignment online | Deployed in production; multi-agent, agent-level simulation | |
| Wang et al. (2018) | Driver | Assignment to a specific order or idle | Location, time, SD features | Trip price | Offline DQN for matching, CFPT network for transfer learning | Single-agent simulation | Single-vehicle problem |
| Tang et al. (2019) and Qin et al. (2020) | Driver | Assignment to a specific order or idle | Location, time, SD features | Trip price | CVNet (deep TD-like algorithm) for learning state values offline + Hungarian method for generating the assignment online | Deployed in production; multi-agent, agent-level simulation | Hierarchical sparse coarse coding, cerebellar embedding of spatial info, Lipschitz regularization on network |
| Holler et al. (2019) | Driver, system | Matching a driver to an order, reposition a driver; matching and repositioning done sequentially | Global info of all drivers and open orders | Trip price, reposition cost | DQN, PPO | Multi-agent, agent-level simulation | Attention mechanism to extract global state info into a context vector |
| Li et al. (2019) | Driver | Assignment to a specific order or idle | Location, time, is_available | Trip price | Mean-field MARL, AC method | Homogeneous vehicles within the same grid cell | The mean action is represented by the peers' destination distribution. |
| Jin et al. (2019) | Worker: hex cell manager: group of hex cells (one layer) | Worker: ranking for match and reposition manager: abstract goal for workers | Number of vehicles, orders, entropy, reposition-guided vehicles, distributions of trip prices and durations in the given hex cell | Manager: total driver income + specifically designed quantity to promote high order response rate worker: intrinsic reward for following the goal generated by manager | Hierarchical MARL | Homogeneous vehicles within the same grid cell | Multiple managers, each manager communicates with multiple workers multi-head attention mechanism for coordination |
| Zhou et al. (2019) | Driver | A trip tuple: (origin cell, destination cell, trip duration, price) | Cell index, number of idle vehicles, orders, distribution of trip destinations in the given hex cell | Trip price | Independent learning with KL divergence regularization | Homogeneous vehicles within the same grid cell | |

(continued on next page)

assigned to consecutive requests since its setting allows a driver to be matched before completing a trip request. [Rong et al. \(2016\)](#) and [Zhou et al. \(2018\)](#) have the direction from which the driver arrives at the current location. For deep RL-based approaches ([Wen et al., 2017](#); [Jiao et al., 2021](#)), richer contextual information, such as SD distributions in the neighborhood, can go into the state.

The learning algorithms are fairly diverse but are all value-based. By estimating the various parameters (e.g., matching probability, passenger destination probability) to compute the transition probabilities, [Rong et al. \(2016\)](#), [Yu et al. \(2019\)](#), [Shou et al. \(2020\)](#) and [Zhou et al. \(2018\)](#) adopt a model-based approach and use value iterations to solve the MDP. [Shou et al. \(2020\)](#) further use inverse RL to learn the unit-distance operational cost. Model-free methods are also common, e.g., Monte Carlo learning ([Verma](#)

Table 2 (continued).

| Paper | Agent | State | Action | Reward | Algorithm | Environment | Notes |
|-----------------------|--------------|---|--|--|--|--|---|
| Ke et al. (2020b) | Trip request | Match or delay | Global: number of idle vehicles, open requests, expected arrival rates of requests and drivers in each cell local: location, cumulative waiting time, expected pick-up distance | Local reward based on the ultimate outcome of matching (whether or not matched or cancelled): trip value, pick-up distance, and match window time Global reward is based on average local reward. Final reward is convex combination of local and global rewards. The rewards have to be updated at the end of the epoch. | DQN, PPO, A2C, ACER with delayed reward. Whole episode trajectories are sampled from replay buffer. | Agent-based simulation with delayed matching feature | |
| Wang et al. (2019) | System | Expected length of current batch | Current nodes in the bipartite graph, current batch size | The sum of the edge weights in the batch | Restricted Q-learning | DiDi GAIA data set | The adaptive batch-based matching has a guarantee on competitive ratio |
| Qin et al. (2021a) | System | Match the current batch or continue to batch (decision made at every time interval) | Number of batched requests and vehicles and estimated arrival rates of demand and supply in each cell | For each time interval, the negative of total matching wait time for all batched requests and the total pick-up wait time saved (by delaying the current batch) | ACER that combines on-policy updates (through a queuing-based simulator) with off-policy updates | Shanghai taxi data | |
| Shi et al. (2019) | Vehicle | Remaining battery level when available, next available time and location, global time | Matching, EV (charging) | Trip price - pick-up cost - charging cost | Similar to CVNet (Tang et al., 2019) | synthetic data | Assumes decomposability of system value into vehicle values |
| Kullman et al. (2022) | System | Global time, new request info, state of each vehicle | Joint matching, reposition, and charging (EV) for each vehicle | Revenue - travel cost | DQN + attention mechanism over vehicles embeddings (similar to Holler et al., 2019) | NYC taxi data with taxi zones | Decision epoch is either a new request arrives or a vehicle becomes idle. |
| Al-Kanj et al. (2020) | System | Supply-demand counts in spatiotemporal discretized space | Joint matching and reposition, EV (charging) | revenue - charging expense | ADP with value function approx. on post-decision states. Value function approx. by hierarchical aggregation. | Simulation with New Jersey ride-hailing trip data | |

et al., 2017), Q-learning (Han et al., 2016; Gao et al., 2018), and DQN (Wen et al., 2017). Jiao et al. (2021) is a hybrid approach in that it performs an action tree search at the online planning stage using estimated matching probabilities and a separately learned state value network. Garg and Ranu (2018) is in a similar spirit by augmenting the multi-arm bandits with Monte Carlo tree search.

The problem formulation most relevant to the ridesharing service provider is system-level vehicle repositioning. Similar to order matching, the ridesharing platform reviews the vehicles' states at fixed time intervals which are significantly longer than those for




Fig. 4. Illustration of the (single-agent) taxi routing problem on a hexagon grid system. The vehicle at its origin position A has the option to reposition to one of the neighboring and current cells. The black arrows represent reposition (idle cruising), and in the two scenarios, the vehicle is matched to a trip request at B and C respectively. The orange arrows represent trip moves, and the orange flags are where the episodes terminate (for long-term horizons). (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

order matching. See Fig. 2 for illustration. Idle vehicles that meet certain conditions, e.g., being idle for sufficiently long time and not in the process of an existing reposition task, are sent reposition recommendations, which specify the desired destinations and the associated time windows. The motivation here is to explicitly modify the current distribution of the available vehicles so that collectively they are better positioned to fulfill more requests more efficiently in the future. Fig. 5 explains the idea with a concrete example. If the vehicles reposition independently (following the orange arrows), they both move to the orange-circled area and there will be a surplus of supply while the demand in the green-circled area will not be served. In contrast, if the vehicles coordinate and the one in the south repositions by the blue arrow, both vehicles will be matched, and all the requests are served.

The agent can be either the platform or a vehicle, latter of which calls for a MARL approach. All the works in this formulation have global SD information (each vehicle and request's status or SD distributions) in the state of the MDP, and a vehicle agent will additionally have its spatiotemporal status in the state. The rewards are mostly the same as in the taxi routing case, except that Mao et al. (2020) consider the monetized passenger waiting time. The actions are all based on grid or taxi zone systems.

The system-agent RL formulation has only been studied very recently, in view of the intractability of the joint action space of all the vehicles (see Table 4). To tackle this challenge of scalability, Feng et al. (2020) decompose the system action into a sequence of atomic actions corresponding to passenger-vehicle matches and vehicle repositions. The MDP encloses a 'sequential decision process' in which all feasible atomic actions are executed to represent one system action, and the MDP advances its state upon complete of the system action. They develop a PPO algorithm for the augmented MDP to determine the sequence of the atomic actions. The system policy in Mao et al. (2020) produces a reposition plan that specifies the number of vehicles to relocate from zone i to j so that the action space is independent from the number of agents (at the expense of additional work at execution). The agent network, trained by a batch AC method, outputs a value for each OD pair, which after normalization gives the percentage of vehicles from each zone to a feasible destination.

The vehicle-agent approaches have to address the coordination issue among the agents. Lin et al. (2018) develop contextual DQN and AC methods, in which coordination is achieved by masking the action space based on the state context and splitting the reward accrued in a grid cell among the multiple agents within the same cell. Oda and Joe-Wong (2018) treat the global state in grid as image input and develop an independent DQN method. They argue that independent learning, equipped with global state information, works quite well compared to an MPC-based approach. The zone structure in Liu et al. (2020) is constructed by clustering a road-connectivity graph. A single vehicle agent is trained with contextual deep RL and generates sequential actions for the vehicles. Zhang et al. (2020b) also train a single DQN agent for all agents, but with global KL distance between the SD distributions similar to Zhou et al. (2019). The DQN agent is put in tandem with QRewriter, another agent with a Q-table value function that converts the output of DQN to an improved action. Shou and Di (2020b) approach the MARL problem with bilevel optimization: The bottom level is a mean-field AC method (Li et al., 2019) with the reward function coming from a platform reward design mechanism, which is tuned by the top level Bayesian optimization. Agent coordination is done by a central module in Chaudhari et al. (2020a), where a vehicle agent executes a mix of independent and coordinated actions. The central module determines the need for coordination based on SD gaps, and explicit coordination is achieved by solving an assignment problem to move vehicles from excess zones to deficit zones.

For joint matching and repositioning optimization, one major challenge is the heterogeneous review cadence. Matching and reposition decisions are typically made asynchronously in practice. To address this issue, Tang et al. (2021) allow the two modules

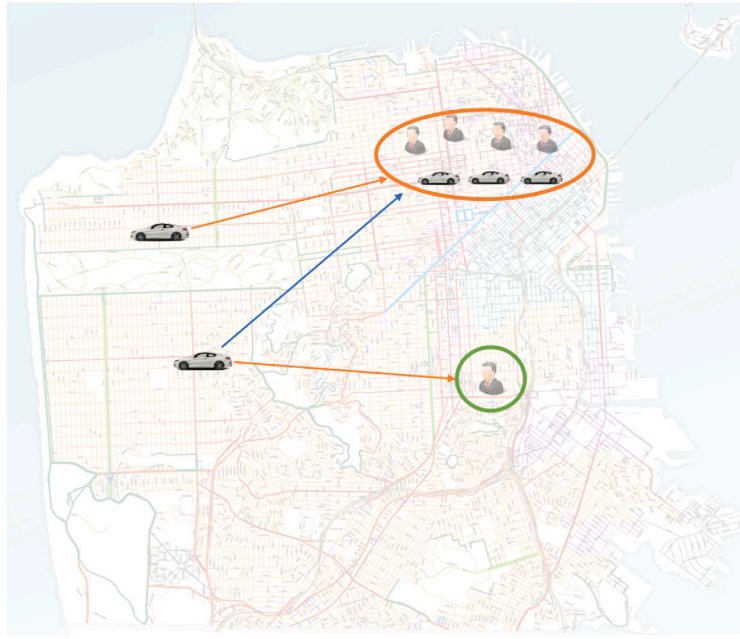


Fig. 5. Illustration of system-level vehicle repositioning. The requests in the orange-circled and green-circled areas appear in the future w.r.t. the time of repositioning. The empty vehicles are existing ones in the orange-circled area. The orange and blue arrows represent potential reposition moves. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

to operate independently but share the same spatiotemporal state value function which is updated online. If the two decisions are formulated into the same problem, the action space can be masked depending on the state (Holler et al., 2019).

Existing RL literature on repositioning often assumes the drivers' full compliance to reposition, i.e., the autonomous vehicle setting. How non-compliance affects the overall performance of a reposition algorithm is a natural question to ask when considering a real-world ridesharing system, in which we expect to see a combination of drivers' independent cruising strategies (Urata et al., 2021; Wong et al., 2014) and system-guided idle cruising behavior. It is also interesting and practically necessary to investigate incentives design and strategies that facilitate the repositioning process. In Zhu et al. (2021), for example, a mean-field MDP is developed for modeling drivers' strategies, and empirical investigations are performed on how spatiotemporal driver incentives affect driver behavior and the system performance.

4.4. Route guidance (navigation)

Routing in this paper refers to low-level navigation decisions on a road network, typically with output of matching and repositioning algorithms as input. The road network, combined with traffic conditions on the links (exhibited as link costs), forms the traffic network which is a non-stationary stochastic network (Mao and Shen, 2018). It is known that standard static shortest-path algorithms do not find the path with minimum expected cost in this case, and the optimal route is not a simple route but a policy (Hall, 1986; Kim et al., 2005). There are two types of set-up for the routing problem, depending on the decision review time. In the first type of set-up, each vehicle on the road network selects a route for a given OD pair from a set of feasible routes. The decision is only reviewed and revised after a trip is completed. Hence, it is called route planning or route choice. When the routes for all the vehicles are planned together, it is equivalent to assigning the vehicles to each link in the network, and hence, the problem is called traffic assignment problem (TAP), which is typically for transportation planning purposes. In the second type of set-up, the routing decision is made at each intersection to select the next outbound road (link) to enter. These are real-time adaptive navigation decisions for vehicles to react to the changing traffic state of the road network. The problem corresponding to this set-up is called dynamic routing, dynamic route choice, or route guidance.

Routing on a road network is a typical multi-agent problem, where the decisions made by one agent has influence on the other agents' performance, simply in that the congestion level of a link depends directly on the number of vehicles passing through that link at a given time and has direct impact on the travel time for all the vehicles on that link within the same time interval. The literature for route planning and TAP often consider the equilibrium property of the algorithms when a population of vehicles adopt them. TAP is typically from a traffic manager's (i.e., system's) perspective. Its goal is to reach system equilibrium (SE, or also often referred to as the system optimum). Some works focus on route planning or TAP from an individual driver's perspective and maximize individual reward. These algorithms try to reach user equilibrium (UE) or Nash equilibrium, under which no agent has the incentive to change its policy because doing so will not achieve higher utility. This is the best that selfish agents can achieve but may not be optimal for the system.

Table 3
Summary of literature for Vehicle Repositioning (taxi routing).

| Paper | Type | Agent | State (in addition to ST info) | Action | Reward | Episode | Algorithm | Coordination | Data |
|----------------------|--------------|---------|---|--|--|--|--|---|--|
| Rong et al. (2016) | Taxi | Driver | Direction from which driver arrives at the current location | Neighboring cells in a grid system | Trip fare | Long-term | Model-based, VI to solve MDP | – | |
| Han et al. (2016) | Taxi | driver | – | Neighboring cells in a grid system | Trip fare - reposition cost | from idle to completion of next trip(s) and being idle again | Q-learning | – | |
| Verma et al. (2017) | Taxi | Driver | – | Neighboring cells in a grid system | Trip fare - reposition cost | Up to the next match | MC learning | – | Taxi log data from Singapore |
| Wen et al. (2017) | Taxi | Driver | SD contextual info | Neighboring cells in a grid system | Saved idle time compared to a counterfactual simulation without reposition. Penalized if no match after reposition. | Up to the next match | DQN with greedy action Compared with MINLP and SAR (simple anticipatory rebalancing) on avg passenger wait time. | – | |
| Garg and Ranu (2018) | Taxi | Driver | Current node on road network | Adjacent node or edge | Idle cruising distance till the next passenger | Up to the next match | MAB + MCTS | – | |
| Gao et al. (2018) | Taxi | Driver | Location, occupied, parking, vacant | Idle cruise to a neighboring cell in the grid system, carrying passenger to destination, waiting | Trip mileage/idle cruising mileage Problem objective: effective driving ratio = total trip mileage/total idle cruising mileage | Long-term (day) | Q-learning | – | Beijing Taxi data 2013 |
| Zhou et al. (2018) | Taxi | Driver | Current road segment, time, previous road segment | Adjacent node or edge | Trip fare | Long-term | Model-based, VI to solve MDP | – | A year of taxi data from a major city in China |
| Yu and Shen (2019) | Taxi | Driver | Current node on road network | Adjacent node or edge | Trip fare - operational cost | long-term (day) | Model-based, VI to solve MDP; use parallel matrix operations to accelerate computation | – | Shanghai taxi trajectory data in the morning of a weekday |
| Shou et al. (2020) | Taxi, system | Driver | Boolean: whether or not assigned to consecutive requests | Neighboring cells in a grid system | Trip fare - operational cost operational cost per unit distance learned through IRL | long-term (day) | Model-based, VI to solve MDP; inverse RL to learn unit distance op cost | Sequentially make decision for each driver. Adjust the matching prob of each driver's MDP, and solve again. | Beijing ride-hailing trajectory data from DiDi over three weekdays |
| Jiao et al. (2021) | Taxi | Vehicle | SD contextual info in the current cell | Neighboring cells in a grid system | Trip fare - reposition cost | Long-term | Offline CVNet + decision-time action search | – | DiDi ride-hailing data |

Table 4
Summary of literature for Vehicle Repositioning (system reposition).

| Paper | Type | Agent | State (in addition to ST info) | Action | Reward | Episode | Algorithm | Coordination | Data |
|--|--------|---------|---|---|--|-----------------|---|--|--|
| Lin et al. (2018) | System | Driver | Global SD contextual features in all cells | Neighboring cells in a grid system | Trip fare, shared when multiple agents are in the same grid cell | Long-term | Contextual DQN, AC | Contextual state features, action space pruned by context | 4 weeks of DiDi data in Chengdu, China |
| Oda and Joe-Wong (2018) | System | Driver | Global SD state discretized into cells, treated as an image | Reachable cells in a grid system within the reposition time cycle | Weighted number of pick-ups - reposition time | Long-term | Independent DQN | – | NYC taxi data |
| Shou and Di (2020b) | System | Vehicle | – | Neighboring cells in a grid system | Trip fare | Long-term | Bilevel optimization: top Bayesian optimization to update reward param, bottom mean-field MARL (AC) | Mean-field MARL | NYC taxi data |
| Jiao et al. (2021) | System | Vehicle | SD contextual info in the current and neighboring cells | Neighboring cells in a grid system | Trip fare - reposition cost | Long-term | Deep SARSA | Stochastic policy through softmax of action values, SD regularization to action values | DiDi ride-hailing data |
| Mao et al. (2020) | System | System | SD info for each zone | Reposition plan: number of repositioned vehicles for each OD pair in a zone map | Monetized passenger waiting time | Long-term | Batch AC | Central decision-making | NYC taxi data |
| Feng et al. (2020) | System | System | Status of every vehicle and request | Atomic action: driver-passenger match or driver-destination match (reposition); system action: sequence of atomic actions | trip fare - operational cost | Long-term (day) | PPO applied to MDP with Sequential Decision Process embedded: global actions decomposed into sequential atomic ones | Central decision-making | DiDi data: 5 regions, 1000 cars, 360 min |
| Liu et al. (2020) | System | Vehicle | Discrete zone structure constructed by clustering a road connectivity graph | Neighboring zones | trip fare | Long-term | Contextual DQN with shared value function | Vehicle actions generated sequentially | real-world taxi data |
| Zhang et al. (2020b) | System | Vehicle | Global SD distributions | Neighboring cells in a grid system | Global KL distance between SD distributions | Long-term | DQN + Q-table in tandem with shared value function | Global state features | DiDi data |
| Chaudhari et al. (2020a) | System | Vehicle | cell in an ST table | Neighboring cells in a grid system, wait | Trip fare - operational cost | Long-term | SARSA-like policy evaluation | Solve an assignment problem of surplus and deficits in terms of SD gap | NYC taxi data |

Value-based RL is by far the most common approach for route planning and TAP aiming to reach UE (see Table 5). In the MDP formulation, the agent is a vehicle (or equivalently, a task) with a given OD pair. The objective is to minimize the total travel time for an individual vehicle (task) (Mainali et al., 2008; Ramos et al., 2018; Zhou et al., 2020b), i.e., the agent is selfish. The immediate reward is the total travel time of a trip for an individual and a particular run. This MDP is stateless, so strictly speaking, it is a multi-arm bandits or contextual bandits problem (Li et al., 2010) if considering time as a contextual feature. The action to take at decision time is to select a route from the set of feasible routes for the associated OD pair (Ramos et al., 2018; Zhou et al., 2020b; Bazzan and Chira, 2015). The value function that governs the route choice decisions represents the long-term expected travel time for the trip identified by the given OD pair. Due to the multi-agent nature, the environment w.r.t. each agent is non-stationary in that the reward function is changing with the policy updates from the other agents. Empirical convergence to UE is demonstrated by Ramos et al. (2018). Zhou et al. (2020b) further develop a Bush–Mosteller RL scheme for MARL and formally establishes its UE convergence property. We also highlight some unique features of the papers. Ramos et al. (2018) consider a different objective from the common and minimizes the driver's regret. To do that, the Q-learning updates are modified using the estimated action regret, which can be computed by local observations and global travel time information communicated by an app. Bazzan and Chira (2015) propose a hybrid method, with Q-learning for individual agents and Genetic Algorithm for reaching system equilibrium, minimizing the average travel time over different trips in the network. This method is thus able to achieve SE. Mainali et al. (2008) adopt Q-iterations with a model set-up similar to that of dynamic routing to be discussed next.

Most applications of RL to routing concern with the *dynamic routing* (DR) problem (see Table 5). The MDP is modeled around a vehicle agent. The basic state information is the traffic state of the current node (i.e., intersection). Some works consider state features of the neighboring nodes (Kim et al., 2005; Mao and Shen, 2018) so that the agent has a broader view of the environment. The action space comprises the set of outbound links (i.e., roads) or adjacent nodes from the current node, so the policy provides a turn-by-turn navigation guidance until the destination is reached. While it is most common to use travel time on a link as the reward function, Tumer et al. (2008) and Grunitzki et al. (2014) stand out by defining a new form called difference reward, which is the difference in average travel time on a link with and without the agent in the system. This applied to only a reward function dependent on the number of agents using the traversed link. In particular, travel distance cannot be used to define a difference reward. Whether solving a specific formulation achieves UE or SE depends on the reward function used. The average travel time on a link is a global reward because it is an aggregate of local rewards (i.e., individual travel times) of all the agents on that link. The difference reward, by definition, is a global reward that also reflects individual effect. If all the agents in the system learn by global reward (Tumer et al., 2008; Grunitzki et al., 2014; Shou and Di, 2020a), then the system is expected to achieve SE. Otherwise, the agents learn by their local rewards, and we will have UE or Nash equilibrium (Kim et al., 2005; Yu et al., 2012; Mao and Shen, 2018; Bazzan and Grunitzki, 2016; Wen et al., 2019).

Most works in the literature adopt Q-learning or its variant as the training algorithm. We report several notable developments. To tackle the sample efficiency issue of online model-free methods, Mao and Shen (2018) propose an offline batch RL approach (fitted Q-iterations) with a tree-based function approximator (Extreme Randomized Trees) that empirically shows good convergence property. Hierarchical methods have also been adopted to address the complexity of a large-scale problem. In Wen et al. (2019), the global road network is divided into sub-networks by differential evolution-based clustering. The top-level network contains only the boundary nodes of the original network. The top-level policy produces the destination node for a sub-network. The sub-level policy provides link-level guidance to reach its sub-destination. Shou and Di (2020b) adopt a bilevel optimization scheme. At the lower level, a mean-field MARL algorithm solves for the dynamic routing problem for the travelers, while at the upper level, a Bayesian optimization module optimizes the control (i.e., reward parameter of the travelers) by the city planner.

4.5. Ride-pooling (carpool)

Ride-pooling optimization typically concerns with matching, repositioning, routing (see e.g., Zheng et al., 2018; Alonso-Mora et al., 2017b,a; Tong et al., 2018). The RL literature has primarily focused on the first two problems. The ride-pooling matching problem differs from that in Section 4.2 in that a combination of multiple passengers, and hence their combined trip, can be matched to a vehicle that may or may not be empty. See stages B and C in Fig. 6 from Alonso-Mora et al. (2017a) for an illustration. The repositioning problem is similar to the ride-hailing case, except that the objective is to optimize some pooling-specific metrics that we define next. The routing problem solves for the sequence of pick-ups and drop-offs given the assigned passengers for a vehicle. The routing problem could also concern with route guidance on the road network. See stage D in Fig. 6.

Many works have multiple objectives and define the reward as a weighted combination of several quantities, with hand-tuned weight parameters. *Passenger wait time* is the duration between the request time and the pick-up time. *Detour delay* is the extra time a passenger spends on the vehicle due to the participation in the ride-pooling. In some cases, these two quantities define the feasibility of a potential pooled trip instead of appearing in the reward (Shah et al., 2020). *Effective trip distance* is the travel distance between the origin and destination of a trip request, should it be fulfilled without ride-pooling. Yu and Shen (2019) consider minimizing passenger wait time, detour delay, and lost demand. Guériau and Dusparic (2018) maximize the number of passengers served. Jindal et al. (2018) maximize the total effective trip distance within an episode, which is just the number of served requests weighted by individual trip distance. Considering a fixed number of requests within an episode (hence fixed maximum effective distance), this metric reflects the efficiency of ride-pooling. Alabbasi et al. (2019), Haliem et al. (2020), Singh et al. (2021) and Haliem et al. (2021) all attempt to minimize the SD mismatch, passenger wait time, reposition time, detour delay, and the number of vehicles used. In addition, Haliem et al. (2020) consider the fleet profit, and Singh et al. (2021) study a more general form of ride-pooling,

Table 5
Summary of literature for Route Guidance.

| Paper | Type | Equilibrium | Agent | State | Action | Reward | Algorithm | Road Network |
|-------------------------|------|--|---------|---|--|--|--|---|
| Mainali et al. (2008) | TAP | UE | Vehicle | Current node (intersection) | An adjacent link | Travel time on the link | Q-iteration. The route is constructed by following the decision at each intersection. | Grid network |
| Bazzan and Chira (2015) | TAP | SE | Vehicle | – | Route from feasible routes for the OD pair | Total travel time on the chosen route | Hybrid method: Q learning for individual agent + Genetic Algorithm for system equilibrium, minimizing avg travel time over different trips. | – |
| Ramos et al. (2018) | TAP | UE (empirical convergence) | Vehicle | – | Route from feasible routes for the OD pair | Total travel time on the chosen route | Q-learning with action-regret updates to minimize driver's total regret | Braess graphs, OW network |
| Zhou et al. (2020b) | TAP | UE (theoretical convergence established) | Vehicle | – | Route from feasible routes for the OD pair | Total travel time on the chosen route | B-M RL scheme, similar to a MARL algo with individual reward. | Nguyen–Dupuis network |
| Kim et al. (2005) | DR | UE | Vehicle | Node, time, binary congestion status vector for all links | An adjacent node | cost accrued by traversing the link | Parameters of MDP estimated from data, MDP solved by value iterations | Southeast Michigan network and traffic data |
| Tumer et al. (2008) | DR | SE | Vehicle | Just a single link | Start time on the link | Difference reward (uplift): with and without the agent in the system | Q-learning | – |
| Yu et al. (2012) | DR | UE | Vehicle | Current node | An adjacent link; action instruction received from the intersection in real time | Travel time on the link | Similar approach to Mainali et al. (2008), but incremental update is done by a step of SARSA. The updates are done in real time and value functions are sync-ed at each intersection, which is an independent traffic management module. | SOUND/4U simulator based on the road network of Kurosaki, Kitakyushu in Japan |

(continued on next page)

where a passenger can hop among different vehicles to complete a trip, with each vehicle completing one leg. They further consider the number of hops and the delay due to hopping.

The state of an agent usually consists of global SD information, similar to that for matching and reposition, but the vehicle status contains key additional information of occupancy and OD's of the passengers on board.

The action space depends on whether the agent is modeled at vehicle level or system level. Existing RL-based works all require that a vehicle drops off all the passengers on board according to a planned route before a new round of pooling. An individual vehicle agent can then match to a feasible group of passengers (in terms of capacity and detour delay) (Jindal et al., 2018), reposition to another location (Alabbasi et al., 2019; Haliem et al., 2020, 2021), or both Guériau and Dusparic (2018). A system-level agent has to make action decisions for the entire fleet together (Yu and Shen, 2019; Shah et al., 2020). The feasible combinations of passengers

Table 5 (continued).

| Paper | Type | Equilibrium | Agent | State | Action | Reward | Algorithm | Road Network |
|---|------|--|----------------|---|------------------|---|---|--|
| Grunitzki et al. (2014) | DR | SE | Vehicle | Current node | An adjacent link | Individual reward: negative travel time experienced by the agent on a link difference reward: as in Tumer et al. (2008) | Applied to a more sophisticated network than aamas08 paper. Results show that DQ-learning outforms IQ-learning. | abstract network topology |
| Mao and Shen (2018) | DR | UE | Vehicle | Current node, time, discrete congestion states vector of the arcs at most 2 steps away. | An adjacent node | Negative travel time experienced by the agent on a link | offline batch RL: fitted Q-iterations with tree-based function approximator (Extreme Randomized Trees). Compared with model-based Q-iterations. | Sioux Falls network |
| Bazzan and Grunitzki (2016) | DR | UE | trip (OD pair) | Current node | An adjacent link | negative travel time experienced by the agent on a link | Independent Q learning compared with successive average method | OW network, Sioux Falls network |
| Wen et al. (2019) | DR | UE | Vehicle | Next approaching node, destination node | An adjacent link | Negative travel time experienced by the agent on a link | tabular Q-learning, global road network clustered into subnetworks by differential evolution-based clustering Top-level network contains only boundary nodes of the original network. Top network policy produces the 'destination' node for a subnetwork. Subnetwork policy provides link-level guidance to reach its sub-destination. | SUMO simulator with various networks in Japan and US |
| Shou and Di (2020b) | DR | SE (avg travel time), UE (travel distance) | Vehicle | Current node, time | An adjacent link | Negative average travel time of all agents or travel distance on a link | Bilevel optimization: Lower level - mean field MARL to solve for dynamic routing for travelers Upper level - Bayesian optimization to optimize controls by city planners | SUMO with Manhattan network |

are typically determined by a separate process based on a pairwise shareability graph or a trip-vehicle graph ([Alonso-Mora et al., 2017a](#)). (See illustration in [Fig. 6](#).)

Papers with vehicle-level policy commonly train a single agent and apply to all the vehicles independently (see e.g., [Haliem et al., 2021](#)). DQN is a convenient choice of training algorithm for this setting. For system-level decision-making, both [Yu and Shen \(2019\)](#) and [Shah et al. \(2020\)](#) employ an ADP approach and consider matching decisions only. [Yu and Shen \(2019\)](#) follow a similar strategy as [Simao et al. \(2009\)](#) and use a linear approximation for the value function. In contrast, [Shah et al. \(2020\)](#) decompose the




Fig. 6. Illustration of the ride-pooling matching process adapted from Alonso-Mora et al. (2017a). Stage A shows the state of the current vehicles and requests. Vehicle 1 has one passenger on board with her destination at the top right-hand corner, while vehicle 2 is empty. The feasible combinations of passengers and the vehicles feasible to serve them are determined at stage B. The corresponding assignment graph is set up and solved at stage C. Stage D shows the resulting routes to fulfill all four new requests as well as the existing trip.

system value function into vehicle-level ones and adopts a neural network for the individual value function, which is updated by mini-batch stochastic gradient descent similar to that in DQN.

It has become increasingly clear that dynamic routing and route planning in the context of ride-pooling require specific attention. In particular, there are two aspects unique to ride-pooling. First, the trips are known only at their request times. Hence, the routes taken by the pooled vehicles (i.e., the sequences of pick-ups and drop-offs) have to be updated dynamically to account for the newly joined passengers. Tong et al. (2018) and Xu et al. (2020) formulate the route planning problem for ride-pooling and develop efficient DP-based route insertion algorithms for carpool. In Section 4.6, we will see that this is also part of the stochastic dynamic vehicle routing problem. Second, within a given route plan, the route taken by a pooled vehicle from an origin to a destination can affect the chance and quality of its future pooling. Hence, dynamic routing (or route choice) between an OD pair can be optimized in that direction, e.g., Yuen et al. (2019) go beyond the shortest-path to make route recommendations for better chance of pooling. Guériau et al. (2020) evaluate the SAMoD system proposed in Guériau and Dusparic (2018) in a microscopic environment based on SUMO with a traffic congestion-aware (non-RL) routing component. We expect to see more RL-based algorithms for the ride-pooling dynamic routing problems.

4.6. Vehicle routing problem (VRP)

VRP has close connection to the ridesharing problem in that variants of VRP could serve as a subroutine of the ride-pooling problem and could even be used to model the entire ridesharing problem itself. The main challenge, in the context of ridesharing, is that new demand (a pair of pick-up and drop-off locations) appears in an online nature and has to be inserted into the existing route dynamically. So reviewing the RL literature for VRP is not only for the completeness of this survey but also essential for one to appreciate the complexity and challenges in tackling ridesharing via RL.

VRP has many variants in its rich literature, so it is important to be clear on their differences and on the variant that each paper claims to solve. The basic setup of a VRP consists of a transportation network $\mathcal{G} := (\mathcal{V}, \mathcal{E}, w)$ and a fleet of K vehicles, where \mathcal{V} is the set of nodes (customer and depot locations), and \mathcal{E} is the set of edges such that $e_{ij} \in \mathcal{E}$ indicates that it is possible to travel from node v_i to node v_j . $w(x_{ij}) := w_{ij}$ is the edge cost, typically distance or travel time. The depot $v_0 \in \mathcal{V}$ is a special node where all the vehicles depart and return at. The vanilla VRP is to find an optimal set of disjoint routes (one for each vehicle) that start and end at the depot, that collectively cover all the nodes, and whose total cost (summing over all the edges in those routes) is minimized. A *traveling salesman problem* (TSP) is a special (simplified) instance of VRP, in which there is no depot, and the fleet consists of a single vehicle. In a *capacitated VRP* (CVRP), each node has a demand quantity to be fulfilled by one of the vehicles. Each vehicle has a limited capacity, and it starts from the depot with a full load of goods to fulfill the demand of the nodes on its route. The total capacities of the fleet is sufficiently large, and all the demand has to be fulfilled. A CVRP with *split delivery* allows the demand of each node to be fulfilled by multiple vehicles. In a VRP with *time windows*, each node has a delivery window within which the node has to be visited or its demand has to be satisfied. If the time window constraints are *soft*, they can be violated with the price of a penalty that contributes to the total cost function. In some variants of the VRP, the goods for the demand of a node (destination) has to be picked up from another designated non-depot location (origin) before being delivered to it. This is known as a VRP with *pick-up and delivery*, also known as the *dial-a-ride problem* (DARP). If the fleet consists of electric vehicles (EV), the set of nodes also include charging stations, and each EV has a limited battery capacity, before the depletion of which the EV has to reach a charging station to recharge. In practical situations, a VRP or variant can be *stochastic* and *dynamic* (SDVRP), i.e., its parameters (e.g., demand and travel time) are uncertain, and the requests are not known at the beginning but are revealed sequentially throughout the problem period.

The connection between VRP and ridesharing exists at both local and fundamental levels. As a subproblem in ride-pooling, the rerouting problem after a new passenger is matched to the vehicle is a TSP with pick-up and delivery (TSPPD), which is one-vehicle single-tour instance of CVRP with pick-up and delivery. At a fundamental level, (multi-vehicle) ride-pooling is a stochastic dynamic multi-vehicle CVRP with pick-up and delivery. Although there are no explicit time windows, cancellation may occur if waiting time is too long. Ride-hailing is also a special case where the vehicles all have unit capacity, and in this case, matching and routing

merge into one single problem. So the ridesharing problem is an SDCVRP, except that repositioning is an intervention strategy not considered in SDCVRP.

The goal of this section is not to provide a complete survey of the VRP literature but rather to point out the representative or unique works that adopt RL to solve VRPs (see Table 6). A recent review of RL-based methods for solving stochastic dynamic VRP can be found in [Hildebrandt, et al. 2021] and a more general one in Ulmer et al. (2020).

Single-vehicle v.s. multi-vehicle problems. CVRP may appear in different forms, and sometimes the subtle differences may not be stated clearly. Most papers solve the single-vehicle problem where there is only one active vehicle at any time. In the capacitated single vehicle problem, the vehicle can make multiple tours (i.e., passing through the depot multiple times) to fulfill all the demand, but the number of tours is not set in advance. In the multi-vehicle case, a fleet of K vehicles are active simultaneously. For static VRPs, if the number of tours in the single-vehicle case is fixed, then it is equivalent to the multi-vehicle counterpart by treating each tour as a separate vehicle. (For problems with time windows, this equivalence can be achieved by resetting the clock every time a new tour starts.) Otherwise, they are not equivalent in general because the number of tours in the optimal solution for the single-vehicle problem may not be N . For dynamic problems where the requests are not all known a-priori, it is not possible to generate the fixed number of tours in sequence, since one cannot insert a new request to a previous tour. In this case, equivalence can only be achieved by keeping each tour on the same clock and updating the routes with the newly appeared requests at each time step. As we will see below, this would render essentially a multi-vehicle algorithm.

The majority of the RL-based methods for VRP models the agent as a vehicle with the system-state visibility. The state thus consists of two types of information: the vehicle state, which includes the vehicle's current location and remaining capacity (for pick-up and delivery, e.g., Ulmer et al., 2020; James et al., 2019; Joe and Lau, 2020) or inventory (for homogeneous goods delivery, e.g., Nazari et al., 2018; Kool et al., 2018; Delarue et al., 2020); the system state, which contains the locations of the customer nodes, the demand at each node, and the unserved customers. For pick-up and delivery problems, the system state instead contains the pick-up and delivery locations of the orders. In the case of EVs, the vehicle state additionally contains the vehicle's battery level, and the system state also includes the locations of the charging stations and the number of vehicles available (not in charging). The action of the agent is to specify the next stop (pick-up/delivery location or charging station in the case EV) to visit for the current vehicle. The sequence of actions form a route for the vehicle. When multiple routes/tours are required, the different routes are separated by the insertion of the depot (Nazari et al., 2018; Duan et al., 2020; Kool et al., 2018; Lin et al., 2021). For (dynamic) multi-vehicle problem where decisions for all the vehicles are made at each time step, the agents would generate their actions sequentially to avoid conflicting actions (James et al., 2019; Zhang et al., 2020a). Since the objective of VRPs is typically to minimize total travel distance, the reward is naturally defined as the negative travel distance. For problems with (soft) time window constraints, the negative penalty for constraint violation is added to the reward.

Typically, these methods adopt an encoder-decoder agent network architecture. The encoder is responsible for encoding part or all of the state information into an embedding vector (or context), which, potentially with additional input state features, is fed into the decoder to generate the action one at a time. Bello et al. (2016) develop a policy network based on the pointer network (Vinyals et al., 2015), which consists of an RNN encoder and an RNN decoder. The major novelty over a sequence-to-sequence architecture is that the decoder uses attention mechanism to attend over the embeddings of the input nodes to generate the probability distribution over the input space, thus eliminating the distance disparity in the output with respect to the input, a feature that is particularly suitable for solving TSP and VRP. To reduce the complexity of the encoder and avoid imposing a sequence on the input state features (e.g., customer locations), which is unnecessary in routing problems, Nazari et al. (2018) modify the pointer network with a non-sequential encoder which simply embeds each individual input node. They incorporate the policy network into an AC method and validate the design on a CVRP with split delivery. A few more recent works have adopted this network structure. James et al. (2019) use structural graph embedding (Struct2Vec) for the encoder, since their agent's state additionally contains a vehicle tour graph. In Lin et al. (2021), the encoder has 1D convolution and graph embedding for the input nodes, followed by an attention layer. Duan et al. (2020) include edge features in the state besides the node features of the transportation network. Their encoder is based on graph convolution network with both node and edge inputs. Another work with significant novelty is Kool et al. (2018), which develops a policy network with a transformer-based encoder and a self-attention-based decoder to use in a PG method (REINFORCE) with the baseline computed from deterministic greedy rollout. This training framework has also been adopted by Zhang et al. (2020a) for multi-vehicle VRP with soft time windows, Lin et al. (2021) for EV VRP with time windows, and Duan et al. (2020), which jointly train an MLP-based binary classifier on edge encoding with the policy network output as labels. They have tested their method on a CVRP with 400 nodes, the largest among the reviewed works.

For SDVRP, Ulmer et al. (2020) argue that it is a more convenient model, which also aligns better with popular approaches to this problem, that the action contains also the route plan information. They define a new variant of MDP, called route-based MDP, in which the state includes the route plan from the last epoch, and the action contains the updated route plan in addition to the next stop to visit. The 'immediate' reward becomes the difference in route value between the old and new plans. Following this line, Joe and Lau (2020) model a system agent whose state includes the cost for the remaining route for each vehicle, and the agent assigns a new request to a vehicle at each decision epoch. The rerouting after matching is solved by simulated annealing for VRP. Under this framework, one only needs to learn an action-value function to generate the matching decisions. The algorithm is tested on a multi-vehicle SDVRP with pick-up/delivery and time windows.⁵ In a somewhat similar spirit but for static CVRP, the MDP action in Delarue et al. (2020) is to generate one route (tour). The value network consists of dense layers and ReLU activation and is representable by mixed-integer linear constraints so that the action can be computed through solving a Prize Collecting TSP by MIP.

⁵ This method can be regarded as one for the matching problem in ride-pooling described in Section 4.5.

Table 6
Summary of literature for VRP.

| Paper | Type | State | Action | Reward | Network | Algorithm | Problem size |
|--------------------------------------|----------------|--|---|---|---|---|--|
| Nazari et al. (2018) | Single-vehicle | Location and demand of each request | The next request to visit | Negative travel distance | Non-sequential encoder for the input with RNN decoder that attends over the input space (Pointer network without an RNN encoder) | AC | Single-vehicle Capacitated VRP with split delivery: one active vehicle at a time |
| Kool et al. (2018) | Single-vehicle | Coordinates and original demand of each node, remaining demand of each node, remaining capacity of the vehicle | Next stop for a given vehicle | Negative travel distance | Transformer encoder with input of coordinates and demand of each node + self-attention-based decoder with additional input of remaining demands and capacity at time t | REINFORCE with greedy rollout baseline | TSP and Capacitated VRP with split delivery, 100 nodes |
| Balaji et al. (2019) | Single-vehicle | Current pickup location, vehicle's location and remaining capacity, orders' locations, statuses, waiting times, and values | Accept an order, pick up an accepted order, wait | Value of delivered order (accept, pickup, deliver) - cost (waiting, traveling, penalty) | Two dense layers NN | APE-X DQN (Horgan et al., 2018), a variant of a DQN that utilizes distributed prioritized experience replay | Stochastic and dynamic CVRP with pick-up/delivery and time windows, 8×8 map, 5 orders 3 pick-up locations |
| James et al. (2019) | Multi-vehicle | System state (available requests, charging station output, vehicles' status(location, battery levels, next stops)), vehicle tour graph | Next stop for a given vehicle; The vehicles generate actions sequentially at each time step. | Expected objective value for a tour | Pointer network for actor, another critic network; Network architecture is similar to NeurIPS-18 paper, but with structural graph embedding (Struct2Vec) for the encoder. | A3C | Multi-vehicle dynamic VRP with pick-up and delivery for EVs: 200 random requests, 100 vehicles |
| Zhang et al. (2020a) | Multi-vehicle | Same as Kool et al. (2018) | Same as Kool et al. (2018) ; The vehicles generate actions sequentially at each time step. | Negative travel distance + negative constraint violation penalty | Same as Kool et al. (2018) | Same as Kool et al. (2018) | Multi-vehicle VRP with soft time windows (no split delivery): 150 nodes, 5 vehicles |
| Ulmer et al. (2020) | Single-vehicle | Vehicle location, time, num of passengers onboard, info of in-process and outstanding requests, route plan from last epoch | The next stop to visit and the new route plan; The paper defines a new variant of MDP called route-based MDP. | Difference in route value between the old route plan and the new one | N.A. | Insert new request s into the current route and use variable neighborhood search to improve the route | SDVRP with pick-up and delivery (DDARP) |

(continued on next page)

4.7. Data sets & environments

The problems in ridesharing are highly practice-oriented, and results from toy data sets or environments may present a very different picture from those in reality. Hence, real-world data sets and realistic simulators backed up by them are instrumental to research in RL algorithms for these problems.

Table 6 (continued).

| Paper | Type | State | Action | Reward | Network | Algorithm | Problem size |
|---------------------------------------|----------------|--|---|---|---|---|---|
| Joe and Lau (2020) | Multi-vehicle | Includes the cost for the remaining route for each vehicle | Matching a new order to a vehicle. Rerouting after matching is done by simulated annealing for VRP | Cost diff between two consecutive decisions | Not specified | NN-based TD learning with experience replay (like in Tang et al. (2019)) to learn the action-value function | Multi-vehicle dynamic VRP with pick-up/delivery and delivery windows: 48 nodes, 2 vehicles, avg 22 orders/day |
| Delarue et al. (2020) | Single-vehicle | The remaining unvisited nodes | To generate one route (tour) through solving a Prize Collecting TSP by MIP | Negative tour distance | Value network consists of dense layers + ReLU activation (representable by mixed-integer linear constraints) | MC policy iteration: rollout N trajectories, fit a new NN | CVRP: 51 nodes |
| Duan et al. (2020) | Single-vehicle | Nodes (location, demand), edges (distance, adjacency) | Generate one node at a time sequentially; The resulting sequence may have multiple depot occurrences for different tours. | Negative travel distance | GCN-based encoder with both node and edge features; GRU-based decoder similar to the pointer network as policy network and MLP-based decoder on the edge encoding as classifier | REINFORCE with greedy rollout baseline (Kool et al., 2018) to train the policy network; Cross-entropy loss to train the binary classifier of route edges with policy output as labels | CVRP: 400 nodes |
| Lin et al. (2021) | Multi-vehicle | For time t , the state of each vertex (location, time window, remaining demand), and global variables (time, battery level of the active vehicle, number of EVs not in charging) | Next stop for the current route; Unlike James et al. (2019) the routes of the vehicles are generated sequentially. Every time the depot appears in the sequence, the system time is reset to 0. | Negative travel distance + negative penalties for constraint violations | Encoder with 1D conv layer and graph embedding for the nodes and attention layer; LSTM-based decoder | REINFORCE with greedy rollout baseline (Kool et al., 2018) | EV with time window and charging. Within the planning horizon, a vehicle can visit the depot only once: C100, S12, EV12 |

The most commonly used data sets are those made available by NYC TLC (Taxi & Limousine Commission) ([TLC, 2020](#)). This large public data repository contains trip records from several different services, Yellow Taxi, Green Taxi, and FHV (For-Hire Vehicle), from 2009 to 2020. The Yellow Taxi data is the most frequently used for various studies. The FHV trip records are submissions from the TLC-licensed bases (e.g., Uber, Lyft) and have a flag indicating pooled trips offered by Uber Pool and Lyft Line. The pick-up and drop-off locations are represented by taxi zones. Manhattan, for example, is divided into 64 zones. There is no driver ID associated with the trip records, so reconstructing historical driver-based trajectories is not possible. An older version of the NYC data set ([Donovan and Work, 2016](#)), however, does include GPS coordinates for pick-up and drop-off locations, and car IDs can be used to track drivers within each year, allowing for more granular and diverse analyses. A similar subset of the NYC FHV data is also available at [Kaggle \(2017\)](#), with GPS coordinates for pick-up and drop-off locations. In addition, travel time data between OD pairs can be obtained through Uber Movement ([Uber, 2021](#)).

Another taxi data set is the San Francisco data set ([Piorkowski et al., 2009](#)), which contains GPS coordinates of approximately 500 taxis collected over 30 days in the San Francisco Bay Area in May 2008. The average time interval between two consecutive location updates is less than 10 s.

A more recent rideshare (Transportation Network Providers, TNP) data set is published by Chicago Data Portal ([Portal, 2020](#)). This data set contains trips, drivers, and vehicles data reported by Transportation Network Providers (TNP, or rideshare companies) in Chicago from 2018. Trip origin and destinations are represented by census tracts. Times are rounded to the nearest 15 min. Fares are rounded to the nearest \$2.50 and tips are rounded to the nearest \$1.00. The driver and vehicle data are not joinable with the trip data.

Developing ridesharing simulators has been a line of research itself. Yao and Bekhor (2021) offer a comprehensive review of recent works on ridesharing simulation models, most of them covering a subset of considerations on the number of passengers, the pre-/post-match passenger cancellation behaviors, and driver acceptance/rejection behaviors. In Yao and Bekhor (2021), a sophisticated event-based simulation framework is proposed to capture all aspects of the behavior modeling. Although the ‘ridesharing’ in their paper is known as the hitch service, where the driver is on her own trip as well, the modeling framework is general and accommodates the ridesharing setting in this survey. Chaudhari et al. (2020a) offer a Gym-compatible, open-source ride-hailing environment (Chaudhari et al., 2020b) for training dispatching and repositioning agents. For large-scale simulation on transport networks, AMoDeus (Ruch et al., 2018) and MATSim (W. Axhausen et al., 2016) are well-established Java-based simulation frameworks that also come with graphical user interfaces and visualization tools. They are of more sophisticated engineering architectures albeit with higher programming bars for extension. The evaluation simulation environment for the KDD Cup 2020 competition is available for public access through the DiDi decision intelligence simulation platform (DiDi, 2021). Although not yet open-sourced, this simulation environment supports both matching and vehicle repositioning tasks and accepts input algorithms through a Python API.

5. Challenges and opportunities

Given the state of the current literature, we discuss a few challenges and opportunities that we feel crucial in advancing RL for ridesharing.

5.1. Ride-pooling

As seen in Section 4.5, the reward function in ride-pooling is often a hand-tuned combination of multiple objectives. It is desirable to have a principled way to determine the best weighting scheme automatically, potentially leveraging inverse RL and multi-objective learning techniques (Zou et al., 2021; Arora and Doshi, 2021) in a similar sense of the ride-hailing case (Zhou et al., 2021). Methods for learning to make matching decisions are still computationally intensive (Shah et al., 2020; Yu and Shen, 2019), in part due to the need to use VRP solver to determine feasible actions (combination of passengers). Moreover, all existing works assume that the action set is pre-determined, and some make only high-level decisions of reposition and serving new passengers or not. A more sophisticated agent may be called for to figure out, for example, how to dynamically determine the desirable passenger combination to match to a vehicle and the routes to take thereafter. Ride-pooling pricing (Ke et al., 2020a), a hard pricing problem itself, is tightly coupled with matching. A joint pricing-matching algorithm for ride-pooling is therefore highly pertinent. As mentioned in Section 4.5, it is also highly anticipated to go beyond using generic routing algorithms and to tailor them to ride-pooling with RL.

5.2. Joint optimization

The rideshare platform is an integrated system, so joint optimization of multiple decision modules leads to better solutions that otherwise unable to realize under separate optimizations, ensuring that different decisions work towards the same goal. RL for joint optimization across multiple modules calls for research on reward function design, state-action representation that facilitates inter-module communication, and the training algorithms. Models and algorithms that allow decentralized execution by the different modules are highly preferred in practice. We have already seen development on RL for joint matching-reposition (Holler et al., 2019; Jin et al., 2019; Tang et al., 2021) and with ride-pooling (Guériau and Dusparic, 2018), pricing-matching (Chen et al., 2019a), and pricing-reposition (Turan et al., 2020). An RL-based method for fully joint optimization of all major modules is highly expected. Meanwhile, this also requires readiness from the rideshare platforms in terms of system architecture and organizational structure.

5.3. Heterogeneous fleet

With the wide adoption of electric vehicles and the emergence of autonomous vehicles, we are facing an increasingly heterogeneous fleet on rideshare platforms. Electric vehicles have limited operational range per their battery capacities. They have to be routed to a charging station when the battery level is low (but sufficiently high to be able to travel to the station). Autonomous vehicles may run within a predefined service geo-fence due to their limited ability (compared to human drivers) to handle complex road situations. For an RL-based approach, a heterogeneous fleet means multiple types of agents with different state and action spaces. The adoption of autonomous vehicles also opens new operational paradigms. Dynamic fleet size inflation (Beirigo et al., 2022), for example, hires idle autonomous vehicles on demand to guarantee service quality contracts in a ridesharing marketplace. Specific studies are required to investigate how to make such a heterogeneous fleet cooperate well to complement each other and maximize the advantage of each type of vehicles to improve overall system efficiency.

5.4. Simulation & Sim2Real

Simulation environments are fundamental infrastructure for successful development of RL methods. Despite those introduced in Section 4.7, simulation continues to be a significant engineering and research challenge. We have rarely seen comparable simulation granularity as that of the environments for traffic management, (e.g., SUMO Lopez et al., 2018, Flow Wu et al., 2017) or autonomous driving (e.g., SMARTS Zhou et al., 2020a, CARLA Dosovitskiy et al., 2017).⁶ The opportunity is an agent-based microscopic simulation environment for ridesharing that accounts for both ride-hailing and carpool, as well as driver and passenger behavior details, e.g., price sensitivity, cancellation behavior, driver entrance/exit behavior. None of the existing public/open-source simulators supports pricing decisions. Those simulators described in the pricing papers all have strong assumptions on passenger and driver price elasticities. A better way might be to learn those behaviors from data through, e.g., generative adversarial imitation learning (Shang et al., 2019) or inverse RL (Mazumdar et al., 2017).

No publicly known ridesharing simulation environment has sufficiently high fidelity to the real world to allow an agent trained entirely in it to deploy directly to production. Several deployed works (Qin et al., 2020; Jiao et al., 2021) in Section 4 have all adopted offline RL for learning the state value functions and online planning. The robotics community has been extensively investigating ways to close the reality gap (Traoré et al., 2019; Mehta et al., 2020). Sim2real transfer algorithms for ridesharing agents are urgently sought after.

5.5. Human behavior

Central to ridesharing platforms are human participants (passengers and drivers).⁷ The impact of human behavior is pervasive in the ridesharing marketplace, e.g., in request conversion, cancellation, idle driver diffusion, driver sign-in and sign-off, rider and driver responses to incentives. Human behavior is inherently stochastic and difficult to model, especially with limited data (in size and features), which introduces errors to optimization and simulation. Compared to traditional approaches from operations research, RL offers potential to better handle these stochasticity issues through its adaptability and data-driven nature.

Unlike cumulative effects induced by spatiotemporal transitions (e.g., matching), human-induced long-term effects from changes in habituation and sentiment on the marketplace are much harder to learn due to the much longer horizon such effects span over. To RL, this is dual challenge and opportunity. The challenge is the long feedback loop and very delayed reward signals, and the opportunities lie in engineering and capturing more refined system state features that capture human behavior characterization better and in designing a richer set of reward signals that facilitate the learning of policies for long-term optimality.

5.6. Non-stationarity

We have seen in Sections 4.2 and 4.3 that RL algorithms deployed to real-world systems generally adopt offline training - once the value function or the policy is deployed, it is not updated until the next deployment. Value functions trained offline using a large amount of historical data are only able to capture recurring patterns resulted from day-on-day SD changes. However, the SD dynamics can be highly non-stationary in that one-time abrupt changes can easily occur due to various events and incidents, e.g., concerts, matches, and even road blocks by traffic accidents. To fully unleash the power of RL, practical mechanisms for real-time on-policy updates of the value function (e.g., Tang et al., 2021; Tong et al., 2021; Eshkevari et al., 2022) is required. In view of the low risk tolerance of production systems in general, sample complexity, computational complexity, and robustness are the key challenges that such methods have to address.

5.7. Business strategies

The research problems in the ridesharing domain are closely associated with how the ridesharing platforms run the operations. Innovation in product and business operations will continue to raise new challenging research problems. There can be multiple alternative product forms to achieve the same goals or address the same challenges, and they inherently define different optimization problems that RL can help tackle.

Surge pricing, for example, is a pricing strategy during the peak hours to address the severe shortage of supply with respect to the surging demand. We have explained its motivation in Section 4.1. While surge pricing is a common practice nowadays, it is not the only strategy that the ridesharing platforms adopt. Passenger requests can be queued if there are no vacant vehicles around to immediately serve the requests (Zhong et al., 2020). The queuing mechanism is perceived in some markets as a more socially acceptable mechanism during the peak hours than surge pricing. Several operational decision questions immediately come up, e.g., how large an area each queue should cover, if the coverage should be dynamically updated, and when the incoming requests should start queuing. These potentially time-varying decisions in a highly stochastic environment are good candidates to be solved for by RL.

Ridesharing platforms often use incentives to stimulate growth on both sides of the marketplace. The forms of incentives are diverse and ever evolving: rider coupons, discounts, target-based challenges, driver bonuses with spatial and temporal constraints,

⁶ Guériau et al. (2020) evaluate the ridesharing algorithms in SUMO, but the environment is not public.

⁷ A partial exception is an autonomous ridesharing platform, where the supply side is powered by autonomous vehicles. However, such services are still prototypical at the time of writing and have very limited coverage.

etc. Each incentive strategy changes the behaviors of a certain segment of the marketplace participants in a certain way, and they inevitably interact with the other marketplace levers, e.g., dynamic pricing (Yang et al., 2020b). The collective effects of the evolving incentives convolute the environment and dynamics of the marketplace, posing significant challenges to RL and other optimization methods. How to represent and capture these factors or explicitly model them in joint optimization is key to tackle these challenges.

Third-party service integrator allows passengers to simultaneously request orders from multiple ride-hailing platforms (Zhou et al., 2022). Service integrators offer the platforms more access to the demand but also bring competition more explicit by displaying the matching information (e.g., trip fare, pick-up distance) side by side. Optimizing pricing and matching policies in a competitive environment with feedback from the service integrator on the competition landscape will be interestingly different from those without a service integrator or in a non-competitive environment. With the added environment complexity, these problems are challenging to solve by traditional methods and could be better tackled by RL.

5.8. General RL

RL provides the necessary tools for the methods reviewed in this survey. Hence, the problems of RL for ridesharing tie closely to the development in RL in general. In the context of ridesharing, we have seen from the literature review above that it is difficult for RL to learn combinatorial actions, e.g., the system matching actions. In the era of deep RL, model interpretability is a long-standing challenge, which hampers investigation of customer experience corner cases. For experience-critical service like ridesharing, policy exploration adds further complication, especially for real-world deployment. In view of these challenges, the future is probably that RL-based and traditional optimization approaches will be complementing each other for a long time. We have seen such combinations in the current literature as Xu et al. (2018) and Qin et al. (2021b) for matching, Chaudhari et al. (2020a) and Jiao et al. (2021) for repositioning, and Delarue et al. (2020) for VRP, that combine RL with combinatorial optimization, mixed-integer programming, and tree search. The breakthroughs of RL that we are seeing in other domains and the continued development of RL methodology for ridesharing certainly make it exciting to anticipate the future landscape.

6. Closing remarks

We have surveyed the RL literature for the core problems in ridesharing: pricing, dispatching, repositioning, routing, ride-pooling, and VRP. We have also discussed some open challenges and future opportunities pertinent to this area.

The ridesharing system is a complex multi-agent system with multiple decision levers. RL offers a powerful modeling vehicle for optimizing this system, but as we have seen from the current literature, challenges remain in tackling complexity in the learning algorithms, the coordination among the agents, and the joint optimization of multiple levers. Along tackling these challenges, we expect that domain knowledge in ridesharing as well as transportation in general will be increasingly instrumental to the successful adoption of RL. As one may have noticed, most of the literature has just appeared in the last four years, and we expect it to continue growing and updating rapidly.

CRedit authorship contribution statement

Zhiwei (Tony) Qin: Conceptualization, Literature survey, Writing – original draft, Writing revision, Project administration.
Hongtu Zhu: Conceptualization, Writing – review & editing. **Jieping Ye:** Conceptualization, Writing – review & editing.

References

- Al-Kanj, L., Nascimento, J., Powell, W.B., 2020. Approximate dynamic programming for planning a ride-hailing system using autonomous fleets of electric vehicles. *European J. Oper. Res.* 284 (3), 1088–1106.
- Alabbasi, A., Ghosh, A., Aggarwal, V., 2019. Deepool: Distributed model-free algorithm for ride-sharing using deep reinforcement learning. *IEEE Trans. Intell. Transp. Syst.* 20 (12), 4714–4727.
- Alonso-Mora, J., Samaranayake, S., Wallar, A., Frazzoli, E., Rus, D., 2017a. On-demand high-capacity ride-sharing via dynamic trip-vehicle assignment. *Proc. Natl. Acad. Sci.* 114 (3), 462–467.
- Alonso-Mora, J., Wallar, A., Rus, D., 2017b. Predictive routing for autonomous mobility-on-demand systems with ride-sharing. In: 2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS). IEEE, pp. 3583–3590.
- Angrist, J.D., Caldwell, S., Hall, J.V., 2021. Uber versus taxi: A driver's eye view. *Am. Econ. J. Appl. Econ.* 13 (3), 272–308.
- Arora, S., Doshi, P., 2021. A survey of inverse reinforcement learning: Challenges, methods and progress. *Artificial Intelligence* 103500.
- Balaji, B., Bell-Masterson, J., Bilgin, E., Damianou, A., Garcia, P.M., Jain, A., Luo, R., Maggior, A., Narayanaswamy, B., Ye, C., 2019. Orl: Reinforcement learning benchmarks for online stochastic optimization problems. *arXiv preprint arXiv:1911.10641*.
- Bazzan, A.L., Chira, C., 2015. A hybrid evolutionary and multiagent reinforcement learning approach to accelerate the computation of traffic assignment. In: *AAMAS*. pp. 1723–1724.
- Bazzan, A.L., Grunitzki, R., 2016. A multiagent reinforcement learning approach to en-route trip building. In: 2016 International Joint Conference on Neural Networks (IJCNN). IEEE, pp. 5288–5295.
- Bei, X., Zhang, S., 2018. Algorithms for trip-vehicle assignment in ride-sharing. In: *Thirty-Second AAAI Conference on Artificial Intelligence*.
- Beirigo, B.A., Negenborn, R.R., Alonso-Mora, J., Schulte, F., 2022. A business class for autonomous mobility-on-demand: Modeling service quality contracts in dynamic ridesharing systems. *Transp. Res. C* 136, 103520.
- Bello, I., Pham, H., Le, Q.V., Norouzi, M., Bengio, S., 2016. Neural combinatorial optimization with reinforcement learning. *arXiv preprint arXiv:1611.09940*.
- Berner, C., Brockman, G., Chan, B., Cheung, V., Debiak, P., Dennison, C., Farhi, D., Fischer, Q., Hashme, S., Hesse, C., et al., 2019. Dota 2 with large scale deep reinforcement learning. *arXiv preprint arXiv:1912.06680*.
- Bertsimas, D., Perakis, G., 2006. Dynamic pricing: A learning approach. In: *Mathematical and Computational Models for Congestion Charging*. Springer, pp. 45–79.

- Bimpikis, K., Candogan, O., Saban, D., 2019. Spatial pricing in ride-sharing networks. *Oper. Res.* 67 (3), 744–769.
- Brown, E., 2020. The ride-hail utopia that got stuck in traffic - WSJ. <https://www.wsj.com/articles/the-ride-hail-utopia-that-got-stuck-in-traffic-11581742802>.
- Chaudhari, H.A., Byers, J.W., Terzi, E., 2020a. Learn to earn: Enabling coordination within a ride hailing fleet. In: *Proceedings of IEEE International Conference on Big Data*.
- Chaudhari, H.A., Byers, J.W., Terzi, E., 2020b. Simulation code for “learn to earn: enabling coordination within a ride-hailing fleet”. <https://transparent-framework.github.io/optimize-ride-sharing-earnings/>.
- Chen, H., Jiao, Y., Qin, Z., Tang, X., Li, H., An, B., Zhu, H., Ye, J., 2019a. InBEDE: Integrating contextual bandit with TD learning for joint pricing and dispatch of ride-hailing platforms. In: *2019 IEEE International Conference on Data Mining (ICDM)*. IEEE, pp. 61–70.
- Chen, Y., Qian, Y., Yao, Y., Wu, Z., Li, R., Zhou, Y., Hu, H., Xu, Y., 2019b. Can sophisticated dispatching strategy acquired by reinforcement learning?—A case study in dynamic courier dispatching system. *arXiv preprint arXiv:1903.02716*.
- Chen, M.K., Sheldon, M., 2016. Dynamic pricing in a labor market: Surge pricing and flexible work on the uber platform. *Ec* 16, 455.
- Chen, C., Yao, F., Mo, D., Zhu, J., Chen, X.M., 2021. Spatial-temporal pricing for ride-sourcing platform with reinforcement learning. *Transp. Res. C* 130, 103272.
- Dantzig, G.B., Ramser, J.H., 1959. The truck dispatching problem. *Manage. Sci.* 6 (1), 80–91.
- Delarue, A., Anderson, R., Tjandraatmadja, C., 2020. Reinforcement learning with combinatorial actions: An application to vehicle routing. *Adv. Neural Inf. Process. Syst.* 33, 609–620.
- DiDi, 2021. Didi decision intelligence simulation platform. <https://outreach.didichuxing.com/Simulation>.
- Donovan, B., Work, D., 2016. New york city taxi trip data (2010–2013). URL: <https://doi.org/10.13012/J8PN93H8>.
- Dosovitskiy, A., Ros, G., Codevilla, F., Lopez, A., Koltun, V., 2017. CARLA: An open urban driving simulator. In: *Conference on Robot Learning*. PMLR, pp. 1–16.
- Duan, L., Zhan, Y., Hu, H., Gong, Y., Wei, J., Zhang, X., Xu, Y., 2020. Efficiently solving the practical vehicle routing problem: a novel joint learning approach. In: *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*. pp. 3054–3063.
- Eshkevari, S.S., Tang, X., Qin, Z., Mei, J., Zhang, C., Meng, Q., Xu, J., 2022. Reinforcement learning in the wild: scalable RL dispatching algorithm deployed in ridehailing marketplace. In: *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*.
- Feng, J., Gluzman, M., Dai, J., 2020. Scalable deep reinforcement learning for ride-hailing. *IEEE Control Syst. Lett.*
- Gao, Y., Jiang, D., Xu, Y., 2018. Optimize taxi driving strategies based on reinforcement learning. *Int. J. Geogr. Inf. Sci.* 32 (8), 1677–1696.
- Garg, N., Ranu, S., 2018. Route recommendations for idle taxi drivers: Find me the shortest route to a customer! In: *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*. pp. 1425–1434.
- Grunitzki, R., de Oliveira Ramos, G., Bazzan, A.L.C., 2014. Individual versus difference rewards on reinforcement learning for route choice. In: *2014 Brazilian Conference on Intelligent Systems*. IEEE, pp. 253–258.
- Guériau, M., Cugurullo, F., Acheampong, R.A., Dusparic, I., 2020. Shared autonomous mobility on demand: A learning-based approach and its performance in the presence of traffic congestion. *IEEE Intell. Transp. Syst. Mag.* 12 (4), 208–218.
- Guériau, M., Dusparic, I., 2018. SAMoD: Shared autonomous mobility-on-demand using decentralized reinforcement learning. In: *2018 21st International Conference on Intelligent Transportation Systems (ITSC)*. IEEE, pp. 1558–1563.
- Haliem, M., Mani, G., Aggarwal, V., Bhargava, B., 2020. A distributed model-free ride-sharing algorithm with pricing using deep reinforcement learning. In: *Computer Science in Cars Symposium*. pp. 1–10.
- Haliem, M., Mani, G., Aggarwal, V., Bhargava, B., 2021. A distributed model-free ride-sharing approach for joint matching, pricing, and dispatching using deep reinforcement learning. *IEEE Trans. Intell. Transp. Syst.* 22 (12), 7931–7942.
- Hall, R.W., 1986. The fastest path through a network with random time-dependent travel times. *Transportation science* 20 (3), 182–188.
- Han, M., Senellart, P., Bressan, S., Wu, H., 2016. Routing an autonomous taxi with reinforcement learning. In: *Proceedings of the 25th ACM International on Conference on Information and Knowledge Management*. pp. 2421–2424.
- Haydari, A., Yilmaz, Y., 2020. Deep reinforcement learning for intelligent transportation systems: A survey. *IEEE Trans. Intell. Transp. Syst.*
- Holler, J., Vuorio, R., Qin, Z., Tang, X., Jiao, Y., Jin, T., Singh, S., Wang, C., Ye, J., 2019. Deep reinforcement learning for multi-driver vehicle dispatching and repositioning problem. In: Wang, J., Shim, K., Wu, X. (Eds.), *2019 IEEE International Conference on Data Mining (ICDM)*. Institute of Electrical and Electronics Engineers, Washington, DC, pp. 1090–1095.
- Horgan, D., Quan, J., Budden, D., Barth-Maroon, G., Hessel, M., Van Hasselt, H., Silver, D., 2018. Distributed prioritized experience replay. *arXiv preprint arXiv:1803.00933*.
- Hu, B., Hu, M., Zhu, H., 2022. Surge pricing and two-sided temporal responses in ride hailing. *Manuf. Serv. Oper. Manage.* 24 (1), 91–109.
- Hu, M., Zhou, Y., 2022. Dynamic type matching. *Manuf. Serv. Oper. Manage.* 24 (1), 125–142.
- James, J., Yu, W., Gu, J., 2019. Online vehicle routing with neural combinatorial optimization and deep reinforcement learning. *IEEE Trans. Intell. Transp. Syst.* 20 (10), 3806–3817.
- Jiao, Y., Tang, X., Qin, Z.T., Li, S., Zhang, F., Zhu, H., Ye, J., 2021. Real-world ride-hailing vehicle repositioning using deep reinforcement learning. *Transp. Res. C* 130, 103289.
- Jin, J., Zhou, M., Zhang, W., Li, M., Guo, Z., Qin, Z., Jiao, Y., Tang, X., Wang, C., Wang, J., et al., 2019. CoRide: Joint order dispatching and fleet management for multi-scale ride-hailing platforms. In: *Proceedings of the 28th ACM International Conference on Information and Knowledge Management*. pp. 1983–1992.
- Jindal, I., Qin, Z.T., Chen, X., Nokleby, M., Ye, J., 2018. Optimizing taxi carpool policies via reinforcement learning and spatio-temporal mining. In: *2018 IEEE International Conference on Big Data (Big Data)*. IEEE, pp. 1417–1426.
- Joe, W., Lau, H.C., 2020. Deep reinforcement learning approach to solve dynamic vehicle routing problem with stochastic customers. In: *Proceedings of the International Conference on Automated Planning and Scheduling*, Vol. 30. pp. 394–402.
- Kaggle, 2017. Uber pickups in New York city - trip data for over 20 million uber (and other for-hire vehicle) trips in NYC. <https://www.kaggle.com/fivethirtyeight/uber-pickups-in-new-york-city>.
- Ke, J., Yang, H., Li, X., Wang, H., Ye, J., 2020a. Pricing and equilibrium in on-demand ride-pooling markets. *Transp. Res. B* 139, 411–431.
- Ke, J., Yang, H., Ye, J., et al., 2020b. Learning to delay in ride-sourcing systems: a multi-agent deep reinforcement learning framework. *IEEE Trans. Knowl. Data Eng.*
- Kim, S., Lewis, M.E., White, C.C., 2005. Optimal vehicle routing with real-time traffic information. *IEEE Trans. Intell. Transp. Syst.* 6 (2), 178–188.
- Kool, W., Van Hoof, H., Welling, M., 2018. Attention, learn to solve routing problems! *arXiv preprint arXiv:1803.08475*.
- Kullman, N.D., Cousineau, M., Goodson, J.C., Mendoza, J.E., 2022. Dynamic ride-hailing with electric vehicles. *Transp. Sci.* 56 (3), 775–794.
- Li, L., Chu, W., Langford, J., Schapire, R.E., 2010. A contextual-bandit approach to personalized news article recommendation. In: *Proceedings of the 19th International Conference on World Wide Web*. pp. 661–670.
- Li, M., Qin, Z., Jiao, Y., Yang, Y., Gong, Z., Wang, J., Wang, C., Wu, G., Ye, J., 2019. Efficient ridesharing order dispatching with mean field multi-agent reinforcement learning. In: *Proceedings of the 2019 World Wide Web Conference on World Wide Web*. International World Wide Web Conferences Steering Committee.
- Lin, L.-J., 1992. Self-improving reactive agents based on reinforcement learning, planning and teaching. *Mach. Learn.* 8 (3–4), 293–321.
- Lin, B., Ghaddar, B., Nathwani, J., 2021. Deep reinforcement learning for the electric vehicle routing problem with time windows. *IEEE Trans. Intell. Transp. Syst.*
- Lin, K., Zhao, R., Xu, Z., Zhou, J., 2018. Efficient large-scale fleet management via multi-agent deep reinforcement learning. In: *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*. pp. 1774–1783.

- Liu, Z., Li, J., Wu, K., 2020. Context-aware taxi dispatching at city-scale using deep reinforcement learning. *IEEE Trans. Intell. Transp. Syst.*
- Lopez, P.A., Behrisch, M., Bieker-Walz, L., Erdmann, J., Flötteröd, Y.-P., Hilbrich, R., Lücken, L., Rummel, J., Wagner, P., Wiefßner, E., 2018. Microscopic traffic simulation using SUMO. In: *The 21st IEEE International Conference on Intelligent Transportation Systems*. IEEE.
- Lowalekar, M., Varakantham, P., Jaillet, P., 2018. Online spatio-temporal matching in stochastic and dynamic domains. *Artificial Intelligence* 261, 71–112.
- Lowe, R., Wu, Y.I., Tamar, A., Harb, J., Pieter Abbeel, O., Mordatch, I., 2017. Multi-agent actor-critic for mixed cooperative-competitive environments. *Adv. Neural Inf. Process. Syst.* 30.
- Lyu, G., Cheung, W.C., Teo, C.-P., Wang, H., 2019. Multi-objective online ride-matching. Available at SSRN 3356823.
- Ma, H., Fang, F., Parkes, D.C., 2020. Spatio-temporal pricing for ridesharing platforms. *ACM SIGecom Exchanges* 18 (2), 53–57.
- Mainali, M.K., Shimada, K., Mabu, S., Hirasawa, K., 2008. Optimal route based on dynamic programming for road networks. *J. Adv. Comput. Intell. Intell. Inform.* 12 (6), 546–553.
- Mao, C., Liu, Y., Shen, Z.-J.M., 2020. Dispatch of autonomous vehicles for taxi services: A deep reinforcement learning approach. *Transp. Res. C* 115, 102626.
- Mao, C., Shen, Z., 2018. A reinforcement learning framework for the adaptive routing problem in stochastic time-dependent network. *Transp. Res. C* 93, 179–197.
- MarketsAndMarkets, 2018. Ride sharing market by type (e-hailing, station-based, car sharing & rental), car sharing (p2p, corporate), service (navigation, payment, information), micro-mobility (bicycle, scooter), vehicle type, and region - global forecast to 2025.
- Mazumdar, E., Ratliff, L.J., Fiez, T., Sastry, S.S., 2017. Gradient-based inverse risk-sensitive reinforcement learning. In: *2017 IEEE 56th Annual Conference on Decision and Control (CDC)*. IEEE, pp. 5796–5801.
- Mehta, B., Deleu, T., Raparthy, S.C., Pal, C.J., Paull, L., 2020. Curriculum in gradient-based meta-reinforcement learning. *arXiv preprint arXiv:2002.07956*.
- Mnih, V., Badia, A.P., Mirza, M., Graves, A., Lillicrap, T., Harley, T., Silver, D., Kavukcuoglu, K., 2016. Asynchronous methods for deep reinforcement learning. In: *International Conference on Machine Learning*. pp. 1928–1937.
- Mnih, V., Kavukcuoglu, K., Silver, D., Rusu, A.A., Veness, J., Bellemare, M.G., Graves, A., Riedmiller, M., Fidjeland, A.K., Ostrovski, G., et al., 2015. Human-level control through deep reinforcement learning. *Nature* 518 (7540), 529–533.
- Nazari, M., Oroojlooy, A., Snyder, L., Takác, M., 2018. Reinforcement learning for solving the vehicle routing problem. In: *Advances in Neural Information Processing Systems*. pp. 9839–9849.
- Ng, A.Y., Russell, S.J., et al., 2000. Algorithms for inverse reinforcement learning. In: *ICML*, Vol. 1. p. 2.
- Oda, T., Joe-Wong, C., 2018. MOVI: A model-free approach to dynamic fleet management. In: *IEEE INFOCOM 2018-IEEE Conference on Computer Communications*. IEEE, pp. 2708–2716.
- Özkan, E., Ward, A.R., 2020. Dynamic matching for real-time ride sharing. *Stoch. Syst.* 10 (1), 29–70.
- Piorkowski, M., Sarafjanovic-Djukic, N., Grossglauser, M., 2009. CRAWDAD dataset epfl/mobility (v. 2009-02-24). Downloaded from <https://crawdad.org/epfl/mobility/20090224>.
- Chicago Data Portal, 2020. Chicago transportation network providers (rideshare) data. <https://data.cityofchicago.org/Transportation/Transportation-Network-Providers-Trips/m6dm-c72p>.
- Powell, W.B., 2007. Approximate Dynamic Programming: Solving the Curses of Dimensionality, Vol. 703. John Wiley & Sons.
- Qin, G., Luo, Q., Yin, Y., Sun, J., Ye, J., 2021a. Optimizing matching time intervals for ride-hailing services using reinforcement learning. *Transp. Res. C* 129, 103239.
- Qin, Z., Tang, X., Jiao, Y., Zhang, F., Xu, Z., Zhu, H., Ye, J., 2020. Ride-hailing order dispatching at DiDi via reinforcement learning. *INFORMS J. Appl. Anal.* 50 (5), 272–286.
- Qin, Z., Zhu, H., Ye, J., 2021b. Reinforcement learning for ridesharing: A survey. In: *Appearing in Proceedings of the IEEE Intelligent Transportation Systems Conference*.
- Raju, C., Narahari, Y., Ravikumar, K., 2003. Reinforcement learning applications in dynamic pricing of retail markets. In: *IEEE International Conference on E-Commerce*, 2003. CEC 2003. IEEE, pp. 339–346.
- Ramos, G.d.O., Bazzan, A.L., da Silva, B.C., 2018. Analysing the impact of travel information for minimising the regret of route choice. *Transp. Res. C* 88, 257–271.
- Rong, H., Zhou, X., Yang, C., Shafiq, Z., Liu, A., 2016. The rich and the poor: A Markov decision process approach to optimizing taxi driver revenue efficiency. In: *Proceedings of the 25th ACM International Conference on Information and Knowledge Management*. pp. 2329–2334.
- Ruch, C., Hörl, S., Frazzoli, E., 2018. Amodeus, a simulation-based testbed for autonomous mobility-on-demand systems. In: *2018 21st International Conference on Intelligent Transportation Systems (ITSC)*. IEEE, pp. 3639–3644.
- Schmoll, S., Schubert, M., 2020. Semi-Markov reinforcement learning for stochastic resource collection. In: *Proceedings of the Twenty-Ninth International Joint Conference on Artificial Intelligence (IJCAI)*.
- Schulman, J., Wolski, F., Dhariwal, P., Radford, A., Klimov, O., 2017. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*.
- Shah, S., Lowalekar, M., Varakantham, P., 2020. Neural approximate dynamic programming for on-demand ride-pooling. In: *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 34. pp. 507–515.
- Shang, W., Li, Q., Qin, Z., Yu, Y., Meng, Y., Ye, J., 2021. Partially observable environment estimation with uplift inference for reinforcement learning based recommendation. *Mach. Learn.* 1–38.
- Shang, W., Yu, Y., Li, Q., Qin, Z., Meng, Y., Ye, J., 2019. Environment reconstruction with hidden confounders for reinforcement learning based recommendation. In: *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*. pp. 566–576.
- Shen, W., He, X., Zhang, C., Ni, Q., Dou, W., Wang, Y., 2020. Auxiliary-task based deep reinforcement learning for participant selection problem in mobile crowdsourcing. In: *Proceedings of the 29th ACM International Conference on Information & Knowledge Management*. pp. 1355–1364.
- Shi, J., Gao, Y., Wang, W., Yu, N., Ioannou, P.A., 2019. Operating electric vehicle fleet for ride-hailing services with reinforcement learning. *IEEE Trans. Intell. Transp. Syst.* 21 (11), 4822–4834.
- Shou, Z., Di, X., 2020a. Multi-agent reinforcement learning for dynamic routing games: A unified paradigm. *arXiv preprint arXiv:2011.10915*.
- Shou, Z., Di, X., 2020b. Reward design for driver repositioning using multi-agent reinforcement learning. *Transp. Res. C* 119, 102738.
- Shou, Z., Di, X., Ye, J., Zhu, H., Zhang, H., Hampshire, R., 2020. Optimal passenger-seeking policies on E-hailing platforms using Markov decision process and imitation learning. *Transp. Res. C* 111, 91–113.
- Silver, D., Hassabis, D., 2016. AlphaGo: Mastering the ancient game of go with machine learning. *Res. Blog* 9.
- Simao, H.P., Day, J., George, A.P., Gifford, T., Nienow, J., Powell, W.B., 2009. An approximate dynamic programming algorithm for large-scale fleet management: A case application. *Transp. Res. C* 13 (2), 178–197.
- Singh, A., Al-Abbasi, A.O., Aggarwal, V., 2021. A distributed model-free algorithm for multi-hop ride-sharing using deep reinforcement learning. *IEEE Trans. Intell. Transp. Syst.*
- Smith, M., 2019. Here's how long you have to wait for an uber or lyft in DC. <https://wtop.com/dc-transit/2019/12/how-long-you-have-to-wait-for-an-uber-or-lyft-in-d-c/>.
- Song, J., Cho, Y.J., Kang, M.H., Hwang, K.Y., 2020. An application of reinforced learning-based dynamic pricing for improvement of ridesharing platform service in Seoul. *Electronics* 9 (11), 1818.
- Sun, H., Wang, H., Wan, Z., 2019. Model and analysis of labor supply for ride-sharing platforms in the presence of sample self-selection and endogeneity. *Transp. Res. B* 125, 76–93.
- Sutton, R.S., 1988. Learning to predict by the methods of temporal differences. *Mach. Learn.* 3 (1), 9–44.

- Sutton, R.S., Barto, A.G., 2018. Reinforcement Learning: An Introduction. MIT Press.
- Tang, X., Qin, Z., Zhang, F., Wang, Z., Xu, Z., Ma, Y., Zhu, H., Ye, J., 2019. A deep value-network based approach for multi-driver order dispatching. In: Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining. pp. 1780–1790.
- Tang, X., Zhang, F., Qin, Z., Wang, Y., Shi, D., Song, B., Tong, Y., Zhu, H., Ye, J., 2021. Value function is all you need: a unified learning framework for ride hailing platforms. In: Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery and Data Mining. In: KDD '21, Association for Computing Machinery, New York, NY, USA, ISBN: 9781450383325, pp. 3605–3615. <http://dx.doi.org/10.1145/3447548.3467096>.
- TLC, 2020. NYC taxi & limousine commission trip record data. <https://www1.nyc.gov/site/tlc/about/tlc-trip-record-data.page>.
- Tong, Y., Shi, D., Xu, Y., Lv, W., Qin, Z., Tang, X., 2021. Combinatorial optimization meets reinforcement learning: Effective taxi order dispatching at large-scale. IEEE Trans. Knowl. Data Eng..
- Tong, Y., Zeng, Y., Zhou, Z., Chen, L., Ye, J., Xu, K., 2018. A unified approach to route planning for shared mobility. Proc. VLDB Endow. 11 (11), 1633.
- Tong, Y., Zhou, Z., Zeng, Y., Chen, L., Shahabi, C., 2020. Spatial crowdsourcing: a survey. VLDB J. 29 (1), 217–250.
- Traoré, R., Caselles-Dupré, H., Lesort, T., Sun, T., Díaz-Rodríguez, N., Filliat, D., 2019. Continual reinforcement learning deployed in real-life using policy distillation and sim2real transfer. arXiv preprint arXiv:1906.04452.
- Tumer, K., Welch, Z.T., Agogino, A., 2008. Aligning social welfare and agent preferences to alleviate traffic congestion. In: Proceedings of the 7th International Joint Conference on Autonomous Agents and Multiagent Systems, Vol. 2. Citeseer, pp. 655–662.
- Turan, B., Pedarsani, R., Alizadeh, M., 2020. Dynamic pricing and fleet management for electric autonomous mobility on demand systems. Transp. Res. C 121, 102829.
- Uber, 2021. Uber movement. <https://movement.uber.com/?lang=en-US>.
- Ulmer, M.W., Goodson, J.C., Mattfeld, D.C., Thomas, B.W., 2020. On modeling stochastic dynamic vehicle routing problems. EURO J. Transp. Logist. 9 (2), 100008.
- Urata, J., Xu, Z., Ke, J., Yin, Y., Wu, G., Yang, H., Ye, J., 2021. Learning ride-sourcing drivers' customer-searching behavior: A dynamic discrete choice approach. Transp. Res. C 130, 103293.
- Verma, T., Varakantham, P., Kraus, S., Lau, H.C., 2017. Augmenting decisions of taxi drivers through reinforcement learning for improving revenues. In: Twenty-Seventh International Conference on Automated Planning and Scheduling.
- Vinyals, O., Fortunato, M., Jaitly, N., 2015. Pointer networks. In: Advances in Neural Information Processing Systems. pp. 2692–2700.
- W. Axhausen, K., Horni, A., Nagel, K., 2016. The Multi-Agent Transport Simulation MATSim. Ubiquity Press.
- Wang, Z., Bapst, V., Heess, N., Mnih, V., Munos, R., Kavukcuoglu, K., de Freitas, N., 2016. Sample efficient actor-critic with experience replay. arXiv preprint arXiv:1611.01224.
- Wang, Z., Qin, Z., Tang, X., Ye, J., Zhu, H., 2018. Deep reinforcement learning with knowledge transfer for online rides order dispatching. In: International Conference on Data Mining. IEEE.
- Wang, Y., Tong, Y., Long, C., Xu, P., Xu, K., Lv, W., 2019. Adaptive dynamic bipartite graph matching: A reinforcement learning approach. In: 2019 IEEE 35th International Conference on Data Engineering (ICDE). IEEE, pp. 1478–1489.
- Wang, H., Yang, H., 2019. Ridesourcing systems: A framework and review. Transp. Res. B 129, 122–155.
- Watkins, C.J., Dayan, P., 1992. Q-learning. Mach. Learn. 8 (3–4), 279–292.
- Wen, F., Wang, X., Xu, X., 2019. Hierarchical sarsa learning based route guidance algorithm. J. Adv. Transp. 2019.
- Wen, J., Zhao, J., Jaillet, P., 2017. Rebalancing shared mobility-on-demand systems: A reinforcement learning approach. In: 2017 IEEE 20th International Conference on Intelligent Transportation Systems (ITSC). IEEE, pp. 220–225.
- Williams, R.J., 1992. Simple statistical gradient-following algorithms for connectionist reinforcement learning. Mach. Learn. 8 (3–4), 229–256.
- Wong, R., Szeto, W., Wong, S., 2014. A cell-based logit-opportunity taxi customer-search model. Transp. Res. C 48, 84–96.
- Wu, T., Joseph, A.D., Russell, S.J., 2016. Automated Pricing Agents in the On-Demand Economy. University of California at Berkeley, Berkeley, CA, USA.
- Wu, C., Kreidieh, A., Parvate, K., Vinitisky, E., Bayen, A.M., 2017. Flow: Architecture and benchmarking for reinforcement learning in traffic control. p. 10, arXiv preprint arXiv:1710.05465.
- Xu, Z., Li, Z., Guan, Q., Zhang, D., Li, Q., Nan, J., Liu, C., Bian, W., Ye, J., 2018. Large-scale order dispatch in on-demand ride-hailing platforms: A learning and planning approach. In: Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining. ACM, pp. 905–913.
- Xu, Y., Tong, Y., Shi, Y., Tao, Q., Xu, K., Li, W., 2020. An efficient insertion operator in dynamic ridesharing services. IEEE Trans. Knowl. Data Eng..
- Yan, C., Zhu, H., Korolko, N., Woodard, D., 2020. Dynamic pricing and matching in ride-hailing platforms. Nav. Res. Logist. 67 (8), 705–724.
- Yang, Y., Luo, R., Li, M., Zhou, M., Zhang, W., Wang, J., 2018. Mean field multi-agent reinforcement learning. In: International Conference on Machine Learning. PMLR, pp. 5571–5580.
- Yang, H., Qin, X., Ke, J., Ye, J., 2020a. Optimizing matching time interval and matching radius in on-demand ride-sourcing markets. Transp. Res. B 131, 84–105.
- Yang, H., Shao, C., Wang, H., Ye, J., 2020b. Integrated reward scheme and surge pricing in a ridesourcing market. Transp. Res. B 134, 126–142.
- Yao, R., Bekhor, S., 2021. A ridesharing simulation platform that considers dynamic supply-demand interactions. arXiv preprint arXiv:2104.13463.
- Yau, K.-L.A., Qadir, J., Khoo, H.L., Ling, M.H., Komisarczuk, P., 2017. A survey on reinforcement learning models and algorithms for traffic signal control. ACM Comput. Surv. 50 (3).
- Yu, X., Gao, S., Hu, X., Park, H., 2019. A Markov decision process approach to vacant taxi routing with e-hailing. Transp. Res. B 121, 114–134.
- Yu, X., Shen, S., 2019. An integrated decomposition and approximate dynamic programming approach for on-demand ride pooling. IEEE Trans. Intell. Transp. Syst. 21 (9), 3811–3820.
- Yu, S., Zhou, J., Li, B., Mabu, S., Hirasawa, K., 2012. Q value-based dynamic programming with SARSA learning for real time route guidance in large scale road networks. In: The 2012 International Joint Conference on Neural Networks (IJCNN). IEEE, pp. 1–7.
- Yuen, C.F., Singh, A.P., Goyal, S., Ranu, S., Bagchi, A., 2019. Beyond shortest paths: Route recommendations for ride-sharing. In: The World Wide Web Conference. pp. 2258–2269.
- Zhang, K., He, F., Zhang, Z., Lin, X., Li, M., 2020a. Multi-vehicle routing problems with soft time windows: A multi-agent reinforcement learning approach. Transp. Res. C 121, 102861.
- Zhang, W., Wang, Q., Li, J., Xu, C., 2020b. Dynamic fleet management with rewriting deep reinforcement learning. IEEE Access 8, 143333–143341.
- Zheng, L., Chen, L., Ye, J., 2018. Order dispatch in price-aware ridesharing. Proc. VLDB Endow. 11 (8), 853–865.
- Zhong, Y., Wan, Z., Shen, Z.-J.M., 2020. Queueing versus surge pricing mechanism: Efficiency, equity, and consumer welfare. In: Equity, and Consumer Welfare (September 24, 2020).
- Zhou, M., Jin, J., Zhang, W., Qin, Z., Jiao, Y., Wang, C., Wu, G., Yu, Y., Ye, J., 2019. Multi-agent reinforcement learning for order-dispatching via order-vehicle distribution matching. In: Proceedings of the 28th ACM International Conference on Information and Knowledge Management. pp. 2645–2653.
- Zhou, F., Lu, C., Tang, X., Zhang, F., Qin, Z., Ye, J., Zhu, H., 2021. Multi-objective distributional reinforcement learning for large-scale order dispatching. In: 2021 IEEE International Conference on Data Mining (ICDM). IEEE, pp. 1541–1546.
- Zhou, M., Luo, J., Villela, J., Yang, Y., Rusu, D., Miao, J., Zhang, W., Alban, M., Fadakar, I., Chen, Z., et al., 2020a. SMARTS: Scalable multi-agent reinforcement learning training school for autonomous driving. arXiv preprint arXiv:2010.09776.
- Zhou, X., Rong, H., Yang, C., Zhang, Q., Khezerlou, A.V., Zheng, H., Shafiq, Z., Liu, A.X., 2018. Optimizing taxi driver profit efficiency: A spatial network-based markov decision process approach. IEEE Trans. Big Data 6 (1), 145–158.
- Zhou, B., Song, Q., Zhao, Z., Liu, T., 2020b. A reinforcement learning scheme for the equilibrium of the in-vehicle route choice problem based on congestion game. Appl. Math. Comput. 371, 124895.
- Zhou, Y., Yang, H., Ke, J., Wang, H., Li, X., 2022. Competition and third-party platform-integration in ride-sourcing markets. Transp. Res. B 159, 76–103.
- Zhu, Z., Ke, J., Wang, H., 2021. A mean-field Markov decision process model for spatial-temporal subsidies in ride-sourcing markets. Transp. Res. B 150, 540–565.
- Zou, F., Yen, G.G., Zhao, C., 2021. Dynamic multiobjective optimization driven by inverse reinforcement learning. Inform. Sci. 575, 468–484.