Data-driven Vehicle Rebalancing with Predictive Prescriptions in the Ride-hailing System

Xiaotong Guo, Qingyi Wang, and Jinhua Zhao

Rebalancing vacant vehicles is one of the most critical strategies in ride-hailing operations. An effective rebalancing strategy can significantly reduce empty miles traveled and reduce customer wait times by better matching supply and demand. While the supply (vehicles) is usually known to the system, future passenger demand is uncertain. There are two ways to handle uncertainty. First, the point-prediction-driven optimization framework involves predicting the future demand and then producing rebalancing decisions based on the predicted demand. Second, the data-driven optimization approaches directly prescribe rebalancing decisions from data. In this study, a predictive prescription framework is introduced to this problem, where the benefits of predictive and data-driven optimization models are combined. Based on a state-of-the-art vehicle rebalancing model, the matching-integrated vehicle rebalancing (MIVR) model, predictive prescriptions are introduced to handle demand uncertainty. Model performances are evaluated using real-world simulations with New York City (NYC) ride-hailing data under four demand scenarios. When demand can be accurately predicted, a point-prediction-driven optimization framework should be adapted. The proposed predictive prescription models achieve shorter customer wait times over the point-prediction-driven optimization models when future demand predictions are not so accurate, and achieve a competitive performance with respect to the cutting-edge robust optimization models.

***Index Terms*—Data-driven Optimization, Predictive Prescriptions, Ride-hailing System, Vehicle Rebalancing**

I. Introduction

Ride-hailing platforms are one of the most essential components of the emerging Mobility-on-Demand (MoD) system, which provides passengers with improved mobility options through a traveler-centric multimodal urban transportation system [1]. With the rapid growth of ride-hailing platforms, such as Uber, Lyft, and DiDi, ride-hailing and ride-sharing services have become increasingly popular all over the world, especially in highly-urbanized regions. In New York City (NYC), ride-hailing platforms transported on average 15 million passengers per month in 2016, which was approximately the same number of trips served by NYC's 43,000 yellow cabs [2]. A recent survey indicates that 36% of American adults have used a ride-hailing platform (Uber or Lyft) in 2018, an increase from 15% in late 2015 [3].

However, ride-hailing platforms face significant challenges with respect to operational efficiency. Despite having algorithmic pricing and matching strategies currently in place, drivers from ride-hailing platforms spend an estimated 40% of the time cruising without passengers in major cities [4]. With technological advances in the field of autonomous driving in the past decade, Autonomous Mobility-on-Demand (AMoD) systems are becoming a reality. With a fleet of autonomous vehicles (AVs), centralized control and planning of vehicles become more vital to efficient operations [5].

One of the major operational decisions critical to the efficient operations of ride-hailing systems is vehicle rebalancing, where vacant vehicles are redistributed proactively to areas with anticipated high demand to reduce the discrepancy between spatial distributions of supply and demand during each time period, therefore reducing customer wait times [6]–[10].

Since the future demand in ride-hailing systems is unknown, performances of rebalancing decisions rely on both the prediction accuracy of future demand and the uncertainty considerations in subsequent optimization. Various machine learning approaches have been developed to produce a point prediction of future demand with high accuracy [11]–[15]. Subsequently, the decisions are made according to either the nominal predicted demand, which is named point-prediction-driven optimization, or an uncertainty set around the prediction, which is termed robust optimization. Robust optimization has been used widely for decision-making under uncertainty and has been applied to the vehicle rebalancing problem in [10]. However, a good demand prediction does not necessarily lead to a good rebalancing decision. In the demand prediction, all errors are considered the same, whereas in the rebalancing problem sending additional vehicles to remote regions due to overestimated future demand would incur a larger cost compared to if the additional vehicles were to be sent to more connected and central regions.

On the other hand, data-driven optimization directly prescribes decisions from data. For example, stochastic optimization has been commonly used for handling problems that require making decisions under uncertainty in Operations Research (OR) [16]. However, standard data-driven optimization approaches, such as Sample Average Approximation (SAA), do not utilize auxiliary information, which leads to an unacceptable waste of good data. To combine ideas from ML and OR while making use of all available observations and information, a data-driven predictive prescriptions framework was proposed to prescribe optimal decisions in decision making under uncertainty [17].

In this paper, a novel data-driven optimization approach, predictive prescription, is introduced into vehicle rebalancing problems to generate better rebalancing decisions against

This manuscript is submitted on Thursday 28th April, 2022.

Xiaotong Guo is with the Department of Civil and Environmental Engineering, Massachusetts Institute of Technology, Cambridge, MA 02139, USA (e-mail: xtguo@mit.edu).

Qingyi Wang is with the Department of Civil and Environmental Engineering, Massachusetts Institute of Technology, Cambridge, MA 02139, USA (e-mail: qingyiw@mit.edu).

Jinhua Zhao is with the Department of Urban Studies and Planning, Massachusetts Institute of Technology, Cambridge, MA 02139, USA (e-mail: jinhua@mit.edu), corresponding author.

demand uncertainty for ride-hailing platforms. The predictive prescriptions are compared with the standard point-prediction-driven optimization framework, stochastic optimization methods, and robust optimization methods. The contributions of this paper can be summarized as follows:

- Introducing the predictive prescription framework into solving the vehicle rebalancing problem in ride-hailing operations.
- Applying the graph convolutional Long Short-Term Memory (LSTM) and the station-based LSTM into predicting the future demand of ride-hailing systems. Simulations results indicate that prediction errors caused by demand underestimation in predictive models can benefit system performances.
- Using real-world simulations to compare model performances of predictive prescription models with point-prediction-driven optimization models under four different demand scenarios. When demand prediction accuracy is low, predictive prescriptions outperform point-prediction-driven optimization in terms of reducing average customer wait times. The edge of data-driven optimization over point-prediction-driven optimization increases when the supply to demand ratio increases. When demand can be predicted accurately, point-prediction-driven optimization is a better approach to adopt.
- Comparing predictive prescriptions with the robust matching-integrated vehicle rebalancing (MIVR) model proposed in [10]. Compared to the robust MIVR model, predictive prescriptions achieve competitive performances without relying on any additional information about the future demand.

The remainder of the paper is structured as follows. Section II reviews the relevant literature in vehicle rebalancing operations, predictive models and data-driven optimization approaches. Section III describes the basic MIVR model and approaches for improving model performances regarding demand uncertainty including predictive methods and data-driven optimization approaches. Data used in this paper is discussed in Section IV. Real-world simulation settings and empirical results are shown in Section V, including performance comparisons between point-prediction-driven optimization models, predictive prescription models, and robust models. Finally, Section VI recaps the main contributions of this work and provides future research directions.

II. Literature Review

A. Vehicle rebalancing

Rebalancing vacant vehicles is a critical operational strategy for ride-hailing platforms in addition to matching customers with drivers [18]. Due to the spatial imbalance of demand and supply in ride-hailing systems, relocating idle vehicles to areas where estimated future demand exceeds vehicle supply could reduce empty miles traveled and customer wait times. An online vehicle rebalancing algorithm developed in [7] led to a 37% reduction in the average customer wait times compared to the scenario where no rebalancing took place.

The vehicle rebalancing problem is first studied in [19], where an adaptive dynamic programming algorithm is proposed for dynamic fleet management with single-period and multi-period travel times.

Since then, various approaches have been proposed to solve the vehicle rebalancing problem in ride-hailing systems. Typical vehicle rebalancing problems discretize the operating region into sub-regions and vacant vehicles are rebalanced between zones by solving a mathematical programming problem. Wen et al. [9] utilized a reinforcement learning approach to address the vehicle rebalancing problem in a shared MoD system. Their proposed method reduced the fleet size by 14% in a real-world simulation in London. Jiao et al. [20] proposed a practical framework based on deep reinforcement learning and decision-time planning for rebalancing vehicles in ride-hailing systems. Braverman et al. [21] designed a fluid-based optimization model to model vehicles in ride-hailing systems. Their proposed method resulted in a higher fraction of passengers served compared to benchmark models. Miao et al. [8] introduced a data-driven distributionally robust vehicle rebalancing model to minimize the worst-case vehicle rebalancing cost, which consists of vehicle rebalancing distance and a service quality function indicating the balanced-ness between supply and demand. Their approach was evaluated with real-world taxi data in NYC and achieved a 30% reduction in idle driving distance on average.

With the advent of autonomous vehicles, vehicle rebalancing problems have been studied extensively for AMoD systems as well in recent years [5]. A fluid model was utilized to model passengers and vehicles, and an optimal rebalancing policy was developed by solving a linear program [22]. A queueing-based theoretical model was also proposed to model the vehicle rebalancing problem in the AMoD system. The algorithm was designed to minimize the total number of rebalancing trips while maintaining vehicle availability [23]. Iglesias et al. [24] proposed a Model Predictive Control (MPC) algorithm to compute rebalancing strategies by leveraging short-term demand forecasts utilizing the LSTM neural networks. Their proposed algorithm significantly reduced the average customer wait time compared to the rebalancing strategy proposed in [22]. In a shared AMoD setting, Tsao et al. [25] proposed an MPC algorithm to optimize routes for both vacant and occupied vehicles.

Besides, decentralized vehicle rebalancing systems were proposed as contingency plans when AVs lost connections with central dispatch systems. Chen et al. [26] proposed a decentralized cooperative cruising method for offline operations of AMoD fleets. Their proposed method shows significant performance improvements compared to strategies with random-selected destinations for rebalancing AVs under different fleet sizes.

Most recently, Guo et al. [10] proposed a MIVR model, introducing driver-customer matching component into the vehicle rebalancing problem to produce better rebalancing decisions. Robust optimization was used to better protect rebalancing decisions against demand uncertainty. Their method could reduce the average customer wait time by 18% compared to approaches proposed in [21] under a real-world simulation

with the NYC ride-hailing data.

One common modeling framework to handle demand uncertainty is to predict and optimize in separate steps, in which a prediction model is built first, followed by an optimization model taking the outputs from the prediction model. Few studies have considered combining prediction and optimization into one framework. Al-kanj et al. [27] studied a sequence of decision problems in a ride-hailing system with autonomous electric vehicles, including vehicle dispatching (matching, rebalancing, EV charging), surge pricing, and fleet size problems. They utilized value functions to represent the spatial and temporal patterns of demand in order to incorporate the downstream impact of a decision made now on the future. The vehicle dispatching problem was modeled as a Markov decision process and addressed with the approximate dynamic programming (ADP) approach. Ramezani and Nourinejad [28] proposed a taxi dispatching model using the model predictive control approach. They incorporated the interrelated impact of normal traffic flows and taxi dynamics when generating dispatching decisions.

In summary, demand prediction and decision-making under uncertainty are two flourishing topics being researched in parallel. In this paper, we will first review literature in demand prediction and decision making under uncertainty separately, and then introduce a data-driven method that optimizes both in one model.

B. Demand Prediction

In recent years, a lot of studies apply deep learning to forecast ride-hailing demand. The state-of-the-art method is the class of Convolutional LSTM (CNN-LSTM) models because of their capacity in capturing the spatiotemporal travel demand patterns. The appropriate variant of the CNN-LSTM model used in travel demand predictions depends on the structure of the problem. Standard CNN is designed to analyze quantities on urban grids, and convolutions are defined with respect to neighboring cells on an imposed artificial grid [11], [12]. Due to the irregular shapes of taxi zones, graph neural networks were used and different types of correlations between spatial entities are defined by adjacency matrices [13]–[15].

All machine learning methods are concerned with selecting the best estimators via Empirical Risk Minimization (ERM), where weights of the network are obtained via gradient-based algorithms such that the empirical average loss is minimized. The loss functions are often standard, differentiable functions: log-likelihood for predicting distributions, cross-entropy for classification, and mean squared error (MSE) for regression. However, this implicitly assumes that the losses for each sample are equally weighted. For example, with an MSE loss function, over(under)-predicting n people yield the same error regardless of the actual demand. However, in downstream applications such as vehicle rebalancing, the actual decision loss of over(under)-predicting a certain amount of demand is highly likely to be different.

C. Decision making under uncertainty in OR

The most widely-used method for decision-making under uncertainty in downstream optimization tasks is stochastic optimization [16]. One traditional method in stochastic optimization is Sample Average Approximation (SAA), where empirical distributions are treated as the true distributions [29]. Another notable approach for decision-making under uncertainty is robust optimization [30], and its data-driven variants [31], where the optimization task considers an uncertainty set around the predicted values and optimizes the worst realization. However, none of the optimization approaches mentioned here utilize auxiliary observations besides the uncertain quantities.

To narrow this gap and combine ML with OR approaches, Bertsimas et al. [17] proposed a predictive prescription framework for decision making under uncertainty where auxiliary observations and data are leveraged to prescribe optimal decisions directly from data in the optimization model. In this paper, the predictive prescription framework is introduced into the vehicle rebalancing problem and it is compared with both point-prediction-driven optimization models with advanced LSTM networks, sample average approximation, and robust optimization models.

III. Methodology

A. Matching-integrated Vehicle Rebalancing Model

In this section, we briefly describe the matching-integrated vehicle rebalancing (MIVR) model proposed by Guo et al. [10], which constitutes the optimization component of proposed data-driven approaches in this paper.

The operational period is divided into Ω identical time intervals indexed by $k = 1, 2, ..., \Omega$, where each time interval has length Δ. The MIVR model is solved in a rolling-horizon manner illustrated in Figure 1, where decision variables are determined repeatedly at the beginning of each time interval. It is worth mentioning that the MIVR model is a forward-looking model incorporating κ future time intervals. When solving the MIVR model at the beginning of time interval k, including the demand during time interval k, κ future time intervals are considered. Only the vehicle rebalancing decisions of the current time interval k will be implemented. Then the vehicle locations are observed and submitted to the MIVR model as inputs for the next time interval.

Additionally, the study area is divided into n sub-regions (zones), where each sub-region i has an estimated demand $r_i^k \geq 0$ at time k. We introduce two sets in this model: i) set of sub-regions $N = \{1, ..., n\}$, and ii) set of time intervals $K = \{1, ..., \kappa\}$.

For each time interval, the MIVR model performs two tasks: i) vehicle rebalancing, which happens at the beginning of each time interval, and ii) driver-customer matching, which is conducted at the end of each time interval.

In the vehicle rebalancing phase, decision variables are represented by $x_{ij}^k \in \mathbb{R}_+$ denoting the number of idle vehicles rebalanced from sub-region i to sub-region j at time k. The number of available vehicles in sub-region i at the end of time interval k is indicated by $S_i^k \in \mathbb{R}_+$. Let d_{ij}^k, w_{ij}^k denote the travel distance and time from sub-region i to sub-region j at time k, respectively, which can be approximated by the distance and travel time between the centroids of two sub-regions. Let $a_{ij}^k \in \{0, 1\}$ denote whether an idle vehicle can be

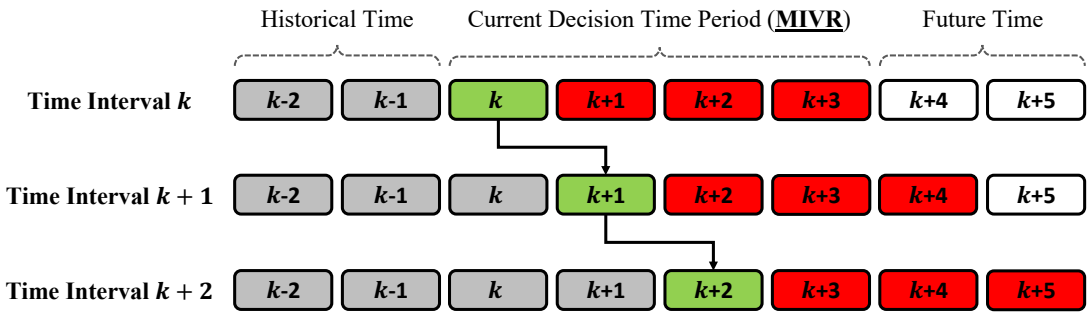

Fig. 1: Example of rolling horizon manner for solving the MIVR model. The MIVR model is solved considering four future time intervals ($\kappa = 4$). Red intervals indicate the look-ahead window while green intervals represent the current decision time intervals whose rebalancing decisions will be implemented.

rebalanced from sub-region i to sub-region j at time k, where $a_{ij}^k = 0$ if rebalancing between sub-regions i, j is feasible at time k. The vehicle rebalancing from sub-region i to sub-region j at time k is feasible if the vehicle can be rebalanced to the destination within the current time interval, i.e., $w_{ij}^k \leq \Delta$. The feasibility constraint for vehicle rebalancing is given by:

$$a_{ij}^k \cdot x_{ij}^k = 0 \quad \forall i, j \in N,\ \forall k \in K. \tag{1}$$

In the MIVR model, since the actual quantity and detailed locations of customers and vehicles are not available, the matching component considers interzonal matchings based on estimated demand. In the matching phase, the decision variables are $y_{ij}^k \in \mathbb{R}_+$ denoting the number of customers in sub-region i matched with vehicles in sub-region j at time k. A maximum pickup time is imposed to guarantee that customers do not experience excessive wait times. Let $\bar{w}$ denote customers' maximum pickup time and parameter $b_{ij}^k \in \{0, 1\}$ denote whether customers in sub-region i can be matched with drivers in sub-region j at time k, where $b_{ij}^k = 0$ indicates a feasible interzonal matching. The matching feasibility constraint is

$$b_{ij}^k \cdot y_{ij}^k = 0 \quad \forall i, j \in N,\ \forall k \in K. \tag{2}$$

The number of unsatisfied requests in sub-region i at time k is represented as $T_i^k \in \mathbb{R}_+$. Then constraints related to the matching phase are:

$$\sum_{j=1}^{n} y_{ji}^k \leq S_i^k \quad \forall i \in N,\ \forall k \in K \tag{3a}$$

$$\sum_{j=1}^{n} y_{ij}^k \leq r_i^k \quad \forall i \in N,\ \forall k \in K \tag{3b}$$

$$T_i^k = r_i^k - \sum_{j=1}^{n} y_{ij}^k \quad \forall i \in N,\ \forall k \in K \tag{3c}$$

Constraints (3a) and (3b) restrict the number of interzonal matchings by the number of available vehicles S_i^k and estimated demand r_i^k. The number of unsatisfied requests is defined as $T_i^k \in \mathbb{R}_+$ by constraints (3c), which is the number of customers who have not been assigned drivers within the current matching phase.

To connect matching and rebalancing phases in the MIVR model, we introduce the following decision variables and parameters:

- $V_i^k \in \mathbb{R}_+$: number of vacant vehicles for sub-region i at the beginning of time interval k.
- $O_i^k \in \mathbb{R}_+$: number of occupied vehicles for sub-region i at the beginning of time interval k.
- $V_i^1, O_i^1, \forall i \in N$: initial vehicle locations.
- $P^k(P_{ij}^k)$: the probability that an occupied vehicle located in sub-region i at time k will be in sub-region j and stay occupied at time $k+1$.
- $Q^k(Q_{ij}^k)$: the probability that an occupied vehicle starting in sub-region i at time k will be in sub-region j and become vacant at time $k+1$.

P^k, Q^k are regional transition matrices describing the movement of occupied vehicles. We approximate them with static matrices estimated from historical data. The approach to estimating these matrices and the limitations of such approximations are discussed in [10].

Then, we specify the following relationships between S_i^k, V_i^k, O_i^k and decision variables x_{ij}^k, y_{ij}^k:

$$\sum_{j=1}^{n} x_{ij}^k \leq V_i^k \quad \forall i \in N,\ \forall k \in K \tag{4a}$$

$$S_i^k = V_i^k + \sum_{j=1}^{n} x_{ji}^k - \sum_{j=1}^{n} x_{ij}^k \quad \forall i \in N,\ \forall k \in K \tag{4b}$$

$$V_i^{k+1} = S_i^k - \sum_{j=1}^{n} y_{ji}^k + \sum_{j=1}^{n} Q_{ji}^k O_j^k \quad \forall i \in N,\ \forall k \in K \setminus \{\kappa\} \tag{4c}$$

$$O_i^{k+1} = \sum_{j=1}^{n} y_{ji}^k + \sum_{j=1}^{n} P_{ji}^k O_j^k \quad \forall i \in N,\ \forall k \in K \setminus \{\kappa\} \tag{4d}$$

Constraint (4a) ensures that the maximum number of vehicles in sub-region i that can be rebalanced to other sub-regions is the number of vacant vehicles at the beginning of time intervals. Constraint (4b) states that available vehicles in sub-region i at time k consist of vacant and rebalanced vehicles. Constraint (4c) shows that vacant vehicles in sub-region i at time $k+1$ are comprised of currently vacant vehicles at time k and currently occupied vehicles that become vacant in the next time interval. Constraint (4d) states that occupied vehicles in

sub-region i at time $k+1$ are comprised of currently vacant vehicles that become occupied in the next interval as well as currently occupied vehicles at time k.

The MIVR model minimizes the number of unsatisfied requests and the total vehicle distance traveled, which consists of vehicle rebalancing distance and vehicle pickup distance. Let γ indicate the penalty (in the unit of VMT[1]) induced by each unsatisfied request, and β defines the relative weighting of rebalancing distance and pickup distance. The MIVR model can be formulated as:

$$\min_{\boldsymbol{x}\in\mathcal{X}} \quad c(\boldsymbol{x};\boldsymbol{r}) = \sum_{k=1}^{\kappa}\sum_{i=1}^{n}\sum_{j=1}^{n} x_{ij}^k d_{ij}^k + \min_{(\boldsymbol{y},\boldsymbol{T})\in\mathcal{L}(\boldsymbol{x},\boldsymbol{r})} \left\{ \gamma \cdot \sum_{k=1}^{\kappa}\sum_{i=1}^{n} T_i^k + \beta \cdot \sum_{k=1}^{\kappa}\sum_{i=1}^{n}\sum_{j=1}^{n} y_{ij}^k d_{ji}^k \right\}, \quad (5)$$

Where

$$\mathcal{L}(\boldsymbol{x},\boldsymbol{r}) = \left\{ (\boldsymbol{y},\boldsymbol{T}) \in \mathbb{R}_+^{n^2\kappa\times n\kappa} : \text{Constraints}(2),(3),(4) \right\},$$
$$\text{and} \quad \mathcal{X} = \left\{ \boldsymbol{x} \in \mathbb{R}_+^{n^2\kappa} : \text{Constraints}(1) \right\}.$$

To simplify the notation, we ignore auxiliary variables $\boldsymbol{S},\boldsymbol{V},\boldsymbol{O}$ in problem (5) and only keep the rebalancing decision vector $\boldsymbol{x}$ and two auxiliary decision vectors $\boldsymbol{y},\boldsymbol{T}$. The demand vector is denoted as $\boldsymbol{r} \in \mathbb{R}_+^{n\kappa}$, which serves as the input parameter of the MIVR model. The MIVR model is a linear programming (LP) problem and can be solved efficiently by off-the-shelf LP solvers, even for large-scale instances (e.g., $n = 500$).

Solving the MIVR model (5) requires the prediction of demand $\boldsymbol{r}$ for future κ time periods. Suppose we are given historical data $(\boldsymbol{z}^i,\boldsymbol{r}^i), i = 1,...,m$, where $\boldsymbol{z}^i \in \mathbb{R}^{n\times\kappa\times p}$ denotes the independent variables with p features, $\boldsymbol{r}^i \in \mathbb{R}^{n\times\kappa}$ is a demand vector which depends upon $\boldsymbol{z}^i$, and m is the number of previous days whose information is provided in the data.

For instance, if we are solving a MIVR model at 9:00 AM today and we would like to predict the future demand $\boldsymbol{r}$ from 9:00 AM to 10:00 AM, we can utilize the historical demand and features between 9:00 AM to 10:00 AM from previous m days, i.e., $\{(\boldsymbol{z}^i,\boldsymbol{r}^i) : \forall i = 1,...,m\}$, to predict the demand today. Meanwhile, we also have access to a feature vector $\boldsymbol{z}$ with exogenous information such as temperature and precipitation for the time period to be predicted.

There are two ways that demand information can be incorporated into the model: point-predictions or data-driven optimization. The former method follows a two-step approach where a point prediction is first produced based on historical observations and auxiliary data independent of the optimization model. Then rebalancing decisions are made according to the point predictions. Data-driven optimization methods, on the other hand, directly prescribe rebalancing decisions from historical observations and auxiliary data.

[1]VMT stands for vehicle miles traveled.

B. Point-Prediction-Driven Optimization

In point-prediction-driven optimization, a predictive model is first developed. Let $f(\cdot)$ represent such a predictive model to predict the unknown demand vector $\boldsymbol{r}$, i.e., $f(\boldsymbol{z}) = \hat{\boldsymbol{r}}$. $f(\cdot)$ can be established based on the data $\{(\boldsymbol{z}_i,\boldsymbol{r}_i), i = 1,...,m\}$ with machine learning methods. The predicted demand $\hat{r}$ is then fed into the MIVR model as $\hat{\boldsymbol{r}}$ to get the rebalancing decisions:

$$\hat{\boldsymbol{x}}^{point-pred} = \operatorname*{argmin}_{\boldsymbol{x}\in\mathcal{X}} c(\boldsymbol{x};\hat{\boldsymbol{r}}). \quad (6)$$

Recent developments on short-term travel demand prediction focus on capturing the spatial-temporal patterns of travel demand using deep learning. The state-of-the-art architecture is the class of Convolutional Long Short Term Memory (Conv-LSTM) networks, where the standard LSTM is extended by having a convolutional structure in both input-to-state and state-to-state transitions [32]. Since sub-regions do not conform to a grid structure, graph convolution proposed by Kipf and Welling [33] is adopted instead of grid convolutions.

Suppose we have L_g graph convolutional layers and the output of the hidden layers is denoted as $H^{(l)}$, $l = 1,...,L_g$, we have the following layer-wise propagation rule:

$$H^{(l+1)} = \sigma(\tilde{D}^{-\frac{1}{2}}\tilde{A}\tilde{D}^{-\frac{1}{2}}H^{(l)}W^{(l)}) \quad (7)$$

where $\sigma(\cdot)$ is an activation function (most commonly ReLU); $\tilde{D}_{ii} = \sum_j \tilde{A}_{ij}$ is the degree matrix; $\tilde{A} = A + I_N$ is the adjacency matrix with added self-connections; and $W^{(l)}$ is the trainable weights of layer l.

Graph convolution layers require upfront access to the global structure of the graph in the form of adjacency matrices (A). In this case, the Euclidean distance between the centroids of the sub-regions is used to relate to neighboring sub-regions.

$$[A_E]_{ij} = \frac{1}{\text{Euclidean Distance}(i,j)} \quad (8)$$

where Euclidean distance is defined as the straight-line distance between the centroids of sub-regions i and j.

In addition to Graph Convolutional LSTM (Graph Conv-LSTM), two LSTM networks without spatial convolution were also constructed as benchmarks. Time series of past demand in each zone are treated as inputs to the model and no spatial correlation between zones is considered. The difference is that in one model, named All Zones LSTM in subsequent discussions, the temporal correlation between different zones is assumed to be the same. Time series from all zones were used to estimate the one-zone model. In the model named Single Zone LSTM, one LSTM is separately trained for each zone. Since All Zones LSTM is a subset of Single Zone LSTM, it is expected that the predictive performance of All Zones LSTM will be the worst among the three models.

C. Data-driven Optimization

Instead of producing a point estimate, there exist data-driven optimization approaches that can prescribe decisions directly from data. First, we consider a simple data-driven approach,

SAA, in this section, which is used as a baseline model. Given a finite sample of data, the SAA approach assumes that the demand vector $\boldsymbol{r}^i$ are drawn uniformly at random from dataset $\{\boldsymbol{r}^i\}_{i=1}^m$. Therefore, the MIVR problem can be written as:

$$\hat{\boldsymbol{x}}^{SAA} = \underset{\boldsymbol{x}\in\mathcal{X}}{\operatorname{argmin}} \frac{1}{m}\sum_{i=1}^{m} c(\boldsymbol{x};\boldsymbol{r}^i). \tag{9}$$

Although SAA accounts for the data uncertainty, it does not utilize any auxiliary information described in $\{\boldsymbol{z}^i\}_{i=1}^m$, which incurs an unacceptable waste of good data. Therefore, we introduce the predictive prescription approach to this problem. Proposed in [17], this framework combines ML and OR techniques and utilizes auxiliary information.

Compared to the traditional SAA approach where only demand vectors $\{\boldsymbol{r}^i\}_{i=1}^m$ are considered for generating rebalancing decisions, the predictive prescription leverages auxiliary observations $\{\boldsymbol{z}^i\}_{i=1}^m$ and solve the following problem:

$$\hat{\boldsymbol{x}}(\boldsymbol{z}) = \underset{\boldsymbol{x}\in\mathcal{X}}{\operatorname{argmin}} \sum_{i=1}^{m} w_i(\boldsymbol{z}) c(\boldsymbol{x};\boldsymbol{r}^i), \tag{10}$$

where $w_i(\boldsymbol{z})$ stands for weight functions derived from historical data $\{(\boldsymbol{z}^i,\boldsymbol{r}^i), i=1,...,m\}$ and current observation $\boldsymbol{z}$. The predictive prescription utilizes machine learning algorithms to generate "smarter" weights compared to identical weights used in the SAA approach.

In this paper, we introduce two machine learning algorithms for generating weights $[w_i(\boldsymbol{z})]_{i=1}^m$. The first algorithm is one of the most commonly used unsupervised learning algorithm, k-nearest-neighbors (KNN). The rebalancing decisions can be generated by solving the following problem:

$$\hat{\boldsymbol{x}}^{KNN}(\boldsymbol{z}) = \underset{\boldsymbol{x}\in\mathcal{X}}{\operatorname{argmin}} \sum_{i\in\mathcal{N}_k(\boldsymbol{z})} c(\boldsymbol{x};\boldsymbol{r}^i), \tag{11}$$

where $\mathcal{N}_k(\boldsymbol{z})$ represents the set of k data points that are closest to $\boldsymbol{z}$, i.e., $\mathcal{N}_k(\boldsymbol{z}) = \{i = 1,...,m : \sum_{j=1}^m \mathbb{1}[\|\boldsymbol{z}-\boldsymbol{z}^i\| \geq \|\boldsymbol{z}-\boldsymbol{z}^j\|] \leq k\}$.

The second algorithm considered in this paper is the optimal regression tree (ORT) proposed in [34], which generates a regression tree with better prediction accuracy than the standard classification and regression tree (CART) approach. The predictive prescription with ORT is formulated as:

$$\hat{\boldsymbol{x}}^{ORT}(\boldsymbol{z}) = \underset{\boldsymbol{x}\in\mathcal{X}}{\operatorname{argmin}} \sum_{i:R(\boldsymbol{z}^i)=R(\boldsymbol{z})} c(\boldsymbol{x};\boldsymbol{r}^i), \tag{12}$$

where $R(\boldsymbol{z})$ corresponds to the leaf of current observation $\boldsymbol{z}$ in the ORT trained on the dataset.

IV. Data Description

The study area is the island of Manhattan in New York City (NYC) and demand data used in this paper is the high-volume ride-hailing trip data [35]. The data includes pickup and drop-off times and locations for all trips made using "high-volume" ride-hailing services, defined as any service that dispatches more than 10,000 trips per day within New York City, including Uber, Lyft, Juno, and Via. We use the data from 20 workdays of June 2019 and the demand is aggregated to 5-minute time intervals.

The sub-regions used in the experiments are "taxi zones" defined within the high-volume ride-hailing trip dataset. There are in total 63 taxi zones on the island of Manhattan. Real travel speed data from June 2019 provided by the Uber Movement database [36] is used for generating interzonal travel times w_{ij}^k. The regional transition probability matrices for occupied vehicles, P^k, and Q^k are generated based on the real travel time and demand data, and details can be found in [10]. Figure 2 shows the mean and standard deviation of daily regional demand in Manhattan. Regions near lower Manhattan have large standard deviations, which imply that accurately predicting demand is not a trivial task when making vehicle rebalancing decisions.

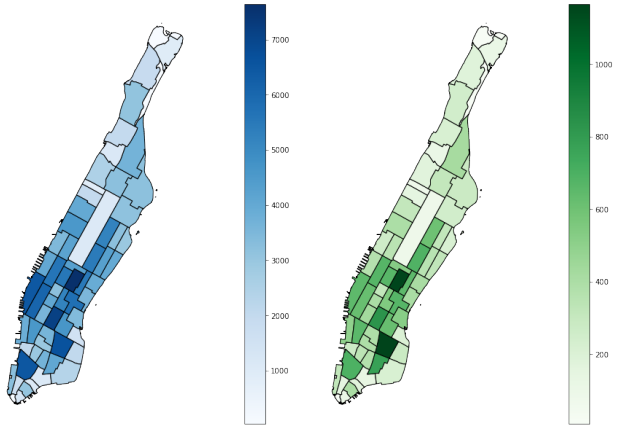

(a) Mean (b) Standard deviation

Fig. 2: Daily demand by zone (trips) in Manhattan.

The auxiliary information used in experiments for both predictive and prescriptive models includes:

- **Weather**: hourly weather data, including air temperature, sensible temperature, precipitation, and snowfall.
- **Point of Interests (POIs)**: number of residential, education, recreational, commercial, and health POIs.
- **Public transit accessibility**: number of subway stations and bus stops.
- **Historical demand**: average demand from previous five-time intervals and historical average demand from m previous days.

Since POI and transit stops/stations are time-independent, they are not used in predictive prescription models.

V. Experimental Results

In this section, we compare model performances of the following approaches: i) point-prediction-driven optimization, ii) SAA, iii) predictive prescriptions and iv) robust MIVR model proposed in [10], as well as two benchmark models: i) optimization with historical average and ii) optimization with true demand under four different demand scenarios. Linear programs in this paper are modeled with open-source Julia [37] package JuMP [38] and solved with Gurobi 9.0 [39] on a 3.0 GHz AMD Threadripper 2970WX Processor with 128 GB Memory.

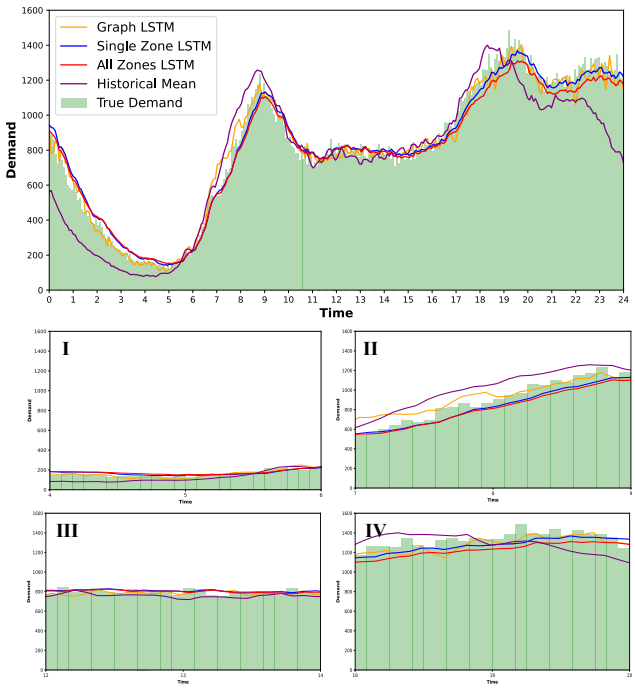

Fig. 3: Demand levels for four different demand scenarios. Green histogram: true demand. Orange line: predicted demand with Graph LSTM. Blue line: predicted demand with Single Zone LSTM. Red line: predicted demand with All Zones LSTM. Purple line: historical average demand.

A. Model Evaluation and Demand Scenarios

To evaluate model performances, we set the last weekday (June 28th, 2019) in our dataset as the test day during which vehicle rebalancing decisions need to be made without knowing the true demand. Data for the previous 19 weekdays are used to construct predictive models and serve as model inputs for data-driven optimization models. Vehicle rebalancing models are evaluated with four different 2-hour demand scenarios which are shown in Figure 3:

- I Morning off-peak scenario (4 - 6): Total demand level is low while point predictions are accurate.
- II Morning peak scenario (7 - 9): Total demand level is high while point predictions are not accurate. Zone-based LSTM underestimates the total demand level.
- III Mid-day off-peak scenario (12 - 14): Total demand level is high while point predictions are accurate.
- IV Evening rush hour scenario (18 - 20): Total demand level is high while point predictions are not accurate. Both zone-based LSTM and graph-based LSTM underestimate the total demand level.

The simulation framework is shown in Figure 4. The input data includes road network for the Manhattan area with shortest path distance and predecessor matrices, distance and travel time matrices between taxi zones, regional transition matrices, demand data, and weights for predictive prescriptions. Parameters used in the simulation are shown in Table I. Fleet size is set to be 2000 vehicles in the simulation. With the setup described above, vehicle and demand locations are

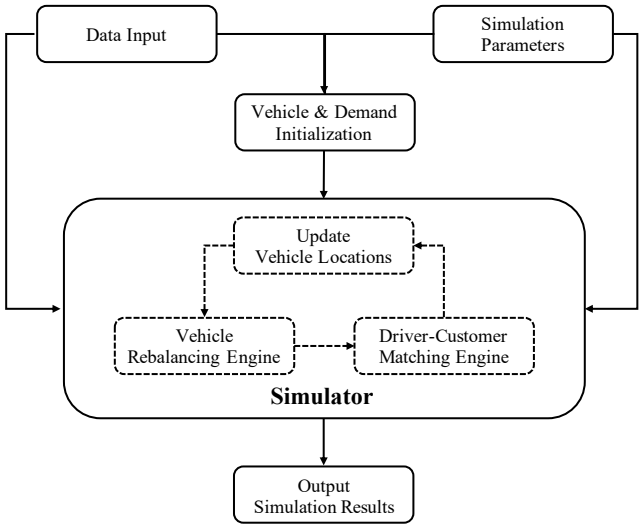

Fig. 4: Simulation framework for evaluating vehicle rebalancing models.

initialized. Vehicles are all available and equally distributed to the taxi zones at the beginning of the simulation. Given that origins and destinations of demand are at the sub-regional level, road nodes within the sub-regions are randomly assigned as origins and destinations for customers in each sub-region.

After initializing vehicle and demand locations, a simulation consisting of a vehicle rebalancing engine and a driver-customer matching engine is run with different rebalancing models. In the simulator, a vehicle rebalancing problem is solved at the start of each time period of length Δ and the vehicle locations are updated before solving vehicle rebalancing problems. A separate driver-customer matching problem is solved at the end of each time period of length δ with available vehicles and realized demand. Details about the driver-customer matching problem can be found in Appendix A. The simulation outputs average customer wait time, unsatisfied customer rate, average non-occupied VMT, and an average number of rebalancing trips for the evaluation of different rebalancing models.

B. Performance of Point Predictions

To ensure that there are enough training samples for neural networks, we utilized additional workday demand data in April and May 2019 in the model training stage. The hyperparameters used in the LSTMs are shown in Table II.

Prediction accuracy for different LSTM models and the benchmark historical average model for the full day are shown in Table III. All machine learning models significantly outperform the historical average. Among the machine learning models, Graph Conv-LSTM has the most representation power, therefore the training error was the smallest. The test set performances for Graph Conv-LSTM and Single Zone LSTM were similar. All Zones LSTM has the worst performance since it does not differentiate demand from different zones.

For the prediction performance under each demand scenario, the MAE is shown in Table IV. For both off-peak demand scenarios (I and III), predictive models have higher prediction

Model Parameter	Explanation	Value
β	Weight parameter for pickup distance	1
γ	Penalty for unsatisfied requests	10^2
Ω	Total number of time intervals in the simulation for each demand scenario	24
Δ	Decision time interval length for vehicle rebalancing problem	300 (seconds)
δ	Batch size for driver-customer matching problem	30 (seconds)
$\bar{w}$	Maximum pickup time	300 (seconds)
$\tilde{w}$	Maximum wait time	300 (seconds)
n	Number of sub-regions	63
κ	Number of look-ahead time intervals when solving MIVR model	6
m	Number of historical data points	19
N_v	Number of vehicles	2000

TABLE I: Model parameters and values.

Hyperparameter	Value
# GCN layers	2
# Units in hidden layers	64
# LSTM layers	1
Weight decay	0.005
Learning rate	0.005

TABLE II: LSTM model setup.

Model	Train MSE	Test MSE	Test MAE
Historical Average	23.18	29.73	3.64
Graph Conv-LSTM	15.63	16.64	2.93
Single Zone LSTM	16.84	16.52	2.93
All Zones LSTM	17.04	18.26	3.04

TABLE III: Prediction performance.

accuracy compared to peak demand scenarios (II and IV). Meanwhile, higher demand leads to higher prediction errors. For peak demand scenarios, zone-based LSTM underestimates the overall demand. Graph-based LSTM only underestimates the overall demand in scenario IV. In the next subsection, we will show that making inaccurate predictions (demand underestimation) could potentially benefit the system's performance.

Model	Scenario I (MAE)	Scenario II (MAE)
Historical Average	1.40	3.64
Graph Conv-LSTM	1.36	3.25
Single Zone LSTM	1.36	3.18
All Zones LSTM	1.40	3.28
Model	Scenario III (MAE)	Scenario IV (MAE)
Historical Average	2.89	4.87
Graph Conv-LSTM	2.81	3.76
Single Zone LSTM	2.85	3.82
All Zones LSTM	2.97	4.20

TABLE IV: Prediction performances under four different demand scenarios.

C. Performance of Different Vehicle Rebalancing Models

In this subsection, we compare the model performances of point-prediction-driven optimization and data-driven optimization, along with two benchmark models: optimization with historical average and optimization with true demand. For predictive models, we constructed a Graph Convolutional LSTM model, a Single Zone LSTM, and an All Zones LSTM model for predicting future demand and generated optimal vehicle rebalancing decisions with point estimations by solving the problem (6). For optimization with the historical average, we used the average demand of m previous workdays as point estimations and solved the problem (6). Similarly, the optimization with true demand utilized the real demand as point estimations and solved the problem (6).

Four data-driven models are considered. The SAA model is included as a benchmark and three predictive prescription models are tested: two KNN models (11) with $k = 5$ and $k = 10$, and an ORT model (12). Weights for m historical days used in predictive prescriptions were generated in the following way. First, a vector $e \in \mathbb{R}^m$ is initialized with m zero values. Then for each unique pair of zones and time intervals, KNN or ORT algorithms were run with m historical data points $\{z_i, i = 1, ..., m\}$ and the current observation z, After that, i-th value in vector e was increased by 1 if z_i is within k nearest neighbors of z or z_i and z belong to the same branch in the constructed ORT. Finally, the weights were generated by normalizing vector e.

Figure 5 shows customer wait times and unsatisfied requests, which are key performance indicators of a ride-hailing system, under four demand scenarios. In each subfigure, colored bars represent the average customer wait time after matching to vehicles while red dotted lines indicate the customer unsatisfaction rate. To better understand how each vehicle rebalancing model works under each demand scenario, the average non-occupied VMT and the average rebalancing trips for each vehicle are shown in Figure 6 and 7. In all four demand scenarios, knowing the true demand leads to the minimum average customer wait time compared to applying any data-driven approaches. The performance comparisons between data-driven optimization and point-prediction-driven optimization vary across different demand scenarios.

In the morning off-peak scenario, the overall demand level is low and all predictive models are more accurate compared to other time periods. Under this demand scenario, point-prediction-driven optimization outperforms data-driven optimization since future demand predictions are very accurate. Figure 5a indicates that data-driven optimization approaches perform even worse than only knowing the historical average demand. On the other hand, data-driven optimization

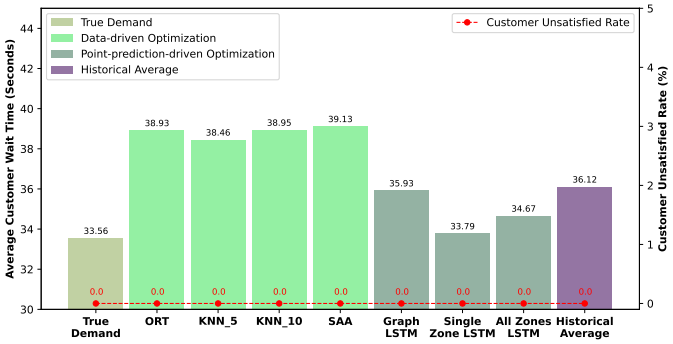

(a) Morning off-peak scenario (4 - 6)

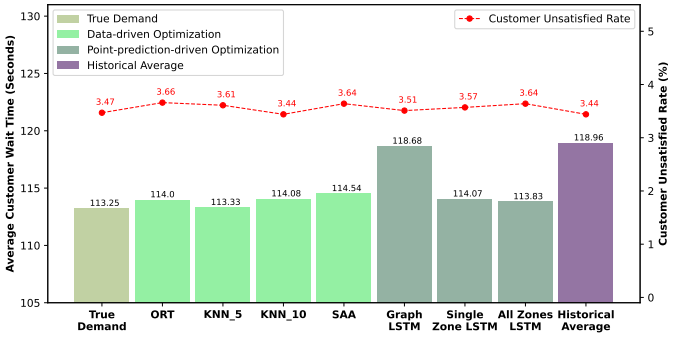

(b) Morning peak scenario (7 - 9)

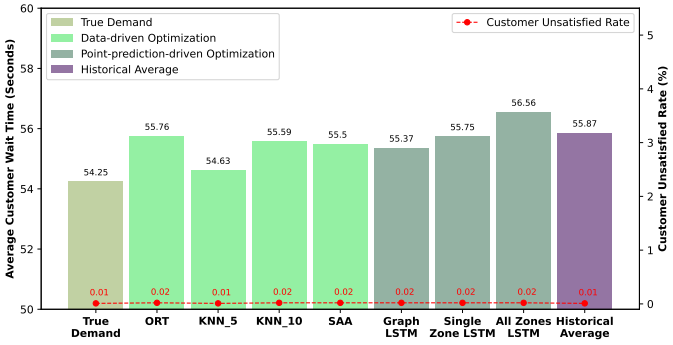

(c) Mid-day off-peak scenario (12 - 14)

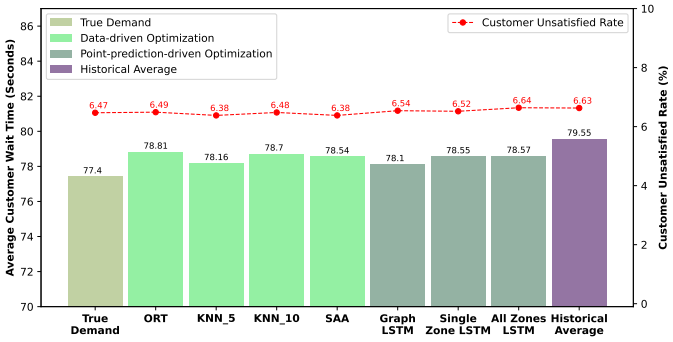

(d) Evening rush hour scenario (18 - 20)

Fig. 5: Customer wait time and unsatisfied rate for different demand scenarios.

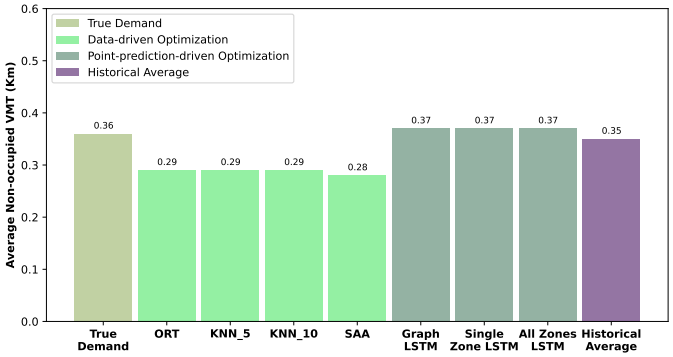

(a) Morning off-peak scenario (4 - 6)

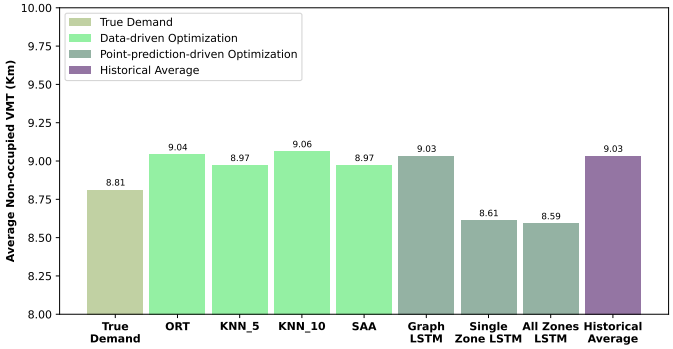

(b) Morning peak scenario (7 - 9)

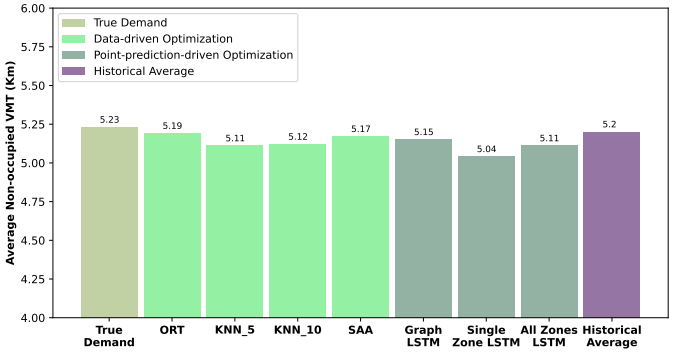

(c) Mid-day off-peak scenario (12 - 14)

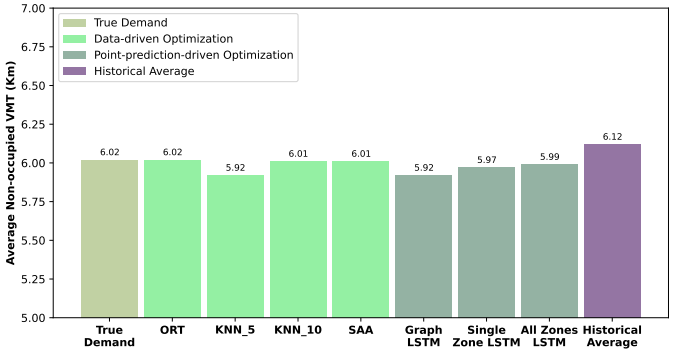

(d) Evening rush hour scenario (18 - 20)

Fig. 6: Average non-occupied VMT for each vehicle under different demand scenarios.

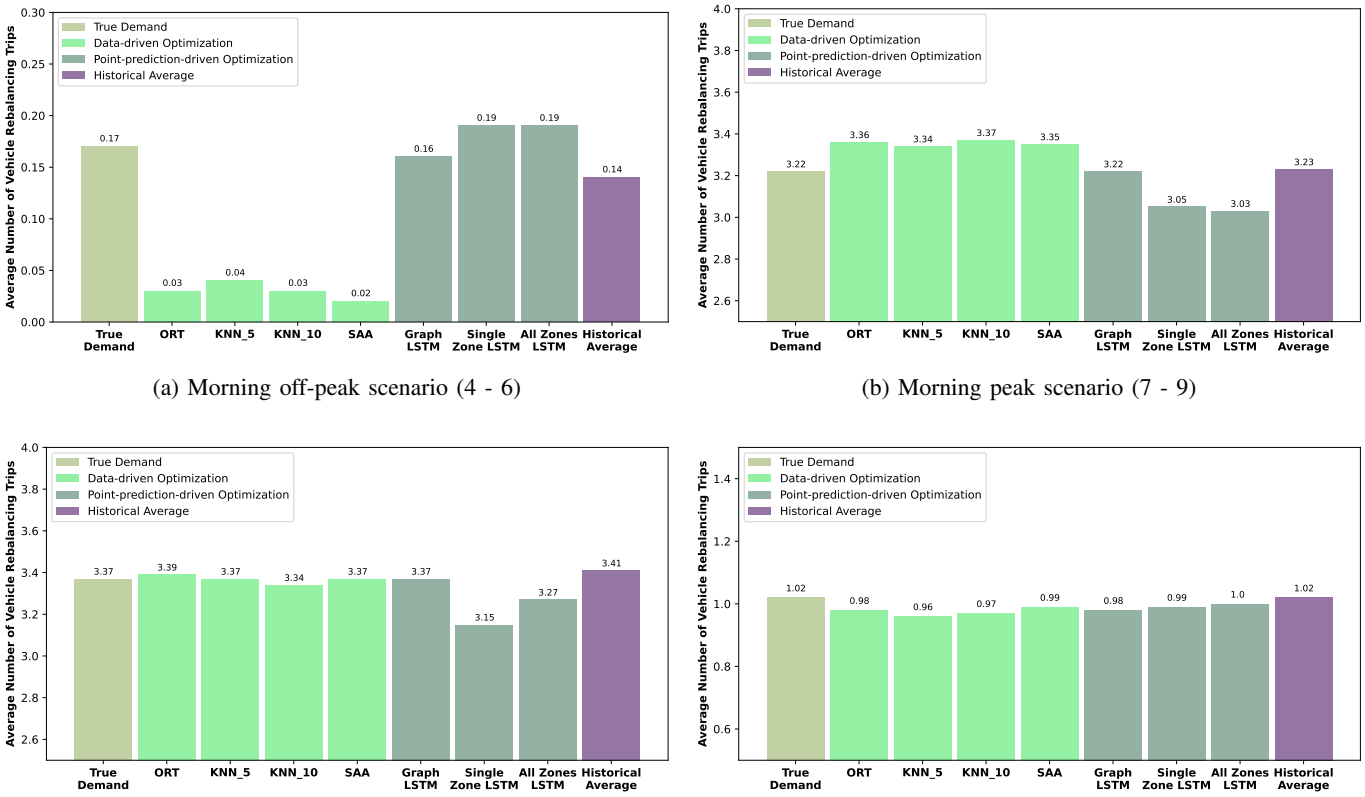

(a) Morning off-peak scenario (4 - 6)

(b) Morning peak scenario (7 - 9)

(c) Mid-day off-peak scenario (12 - 14)

(d) Evening rush hour scenario (18 - 20)

Fig. 7: Average number of rebalancing trips made for each vehicle under different demand scenarios.

approaches conduct much fewer vehicle rebalancing trips according to Figure 7a. When combining with the average non-occupied VMT for each vehicle shown in Figure 6a, we know that data-driven optimization approaches distribute fewer idle vehicles with longer distances. The poor performances of data-driven optimization models imply that several days with low demand levels are deemed more relevant by the model. To summarize, when demand can be accurately predicted, point-prediction-driven optimization should be used.

For the morning peak scenario, the overall demand level is high and predictive models have large prediction errors. Figure 5b shows the average customer wait time and customer unsatisfied rate for each model. The customer unsatisfaction rate is fairly close across all different models. Under this demand scenario, data-driven optimization models perform better overall compared to point-prediction-driven optimization models. All four data-driven optimization models achieve competitive performances with respect to the optimal case in which true demand is known. For predictive models, the graph-based LSTM has the worst performance while two zone-based LSTMs have competitive performances compared to data-driven optimization models.

It is worth mentioning that Graph LSTM has better prediction accuracy than All Zones LSTM though it has a worse model performance. The main reason for zone-based LSTMs to have satisfying performances is that they underestimate future demand, which is shown in Figure 3. Also from Figure 6b and 7b, less rebalancing trips and lower non-occupied VMT imply the demand underestimation by zone-based LSTMs. The "conservativeness" brought by the underestimation leads to better system performances given high volatility in ride-hailing demand. Being conservative is also the key reason for the robust MIVR model proposed in [10] to have satisfying performances. The simulation results justify that a better demand prediction does not necessarily lead to a better rebalancing decision. Meanwhile, underestimation is a more desirable prediction error to make than overestimation when predicting future demand for the purpose of distributing vacant vehicles.

As for the mid-day demand scenario, the overall demand is at a medium level while predictive models are more accurate than the two peak demand scenarios. Under this demand scenario, data-driven optimization models perform better overall compared to point-prediction-driven optimization models based on Figure 5c. Figure 6c and 7c show that they conduct similar number of rebalancing trips with similar distance. All Zones LSTM performs worse than the model knowing the historical average since it has a worse prediction accuracy. When the demand prediction is not accurate enough, data-driven optimization has a close edge over point-prediction-driven optimization.

Under the evening rush hour scenario, the overall demand level is high and predictive models have the worst performances compared to the other three demand scenarios. In this demand scenario, data-driven optimization and point-prediction-driven optimization have similar performances re-

garding the average customer wait time and customer unsatisfaction rate according to Figure 5d. There are limited idle vehicles that can be rebalanced due to a high demand level in the evening rush hour scenario. Figure 7d indicates that the average number of rebalancing trips performed for each vehicle is nearly 1, while the number is over 3 for scenarios II and III with high demand levels. Although data-driven optimization performs better when demand predictions are not accurate, the limited number of idle vehicles leaves no space for data-driven optimization to improve system performances by proactively balancing demand and supply.

Within four data-driven optimization models, predictive prescription with KNN ($k = 5$) performs better than the other three methods by having lower average customer wait times across four demand scenarios. Meanwhile, for scenarios where the demand level is high (morning peak, mid-day off-peak, and evening rush hour), predictive prescription with KNN ($k = 5$) utilizes the minimum VMT over rebalancing idle vehicles. This performance superiority implies that sparsity is an ideal property when applying data-driven optimization. Compared to the predictive prescription with KNN-5, the other three models incorporate more historical demand scenarios, which could diminish the system performance if some demand scenarios are significantly different from the future demand scenario over which rebalancing decisions are made.

To summarize, there are two factors to consider when choosing vehicle rebalancing models: i) supply to demand ratio, and ii) demand prediction accuracy. When the demand can be accurately predicted, point-prediction-driven optimization models perform the best. When the demand is hard to predict (for example, during rush hour), data-driven optimization models perform the best. System performances can be further improved if the supply to demand ratio is higher, where more idle vehicles are available to be rebalanced. Compared to the standard data-driven optimization approach, SAA, predictive prescriptions perform better by leveraging auxiliary information. On the other hand, when demand cannot be accurately predicted, system performances can benefit from underestimation, so fewer unnecessary rebalancing trips are made. However, predictive models tend to aim for "unbiasedness", where the amount of overestimation and underestimation is the same.

D. Comparison with the Robust MIVR Model

In this subsection, we compared the best performing data-driven optimization model, prescriptive prescription with KNN-5, with the robust MIVR model proposed in [10] under the morning peak scenario. We evaluated performances of the robust MIVR model under multiple uncertain scenarios defined by uncertain parameters ρ and Γ via the simulation described in section V-A. Parameters ρ and Γ are parameters defining the size of uncertainty set in the robust MIVR model, and details can be found in Appendix B.

Figure 8 shows the percentage reduction of average customer wait time for the robust MIVR model compared to predictive prescription with KNN-5. Each cell indicates an uncertain scenario (defined by parameters ρ and Γ) in the robust MIVR model. Larger values of uncertain parameters ρ and Γ lead to more conservative rebalancing decisions (since higher demand uncertainty is considered in the model). It is worth mentioning that the uncertain parameter ρ significantly influences the downstream matching performances, while the effect of uncertain parameter Γ is marginal.

In general, the predictive prescription with KNN-5 outperforms the robust MIVR model regarding the average customer wait time. The robust MIVR model could achieve a similar average customer wait time when larger demand uncertainty is considered. On the other hand, the robust MIVR model can satisfy more customers compared to the predictive prescription with KNN-5, which is shown in Figure 9. More customers can be satisfied when considering lower levels of demand uncertainty. The additional customers served by the robust MIVR model are "hard" customers which require longer pickup distances, hence longer wait times. To summarize, predictive prescriptions can reduce the average customer wait time compared to the robust MIVR model. However, a small proportion of customers will not be satisfied, which is likely the reason behind reduced customer wait times.

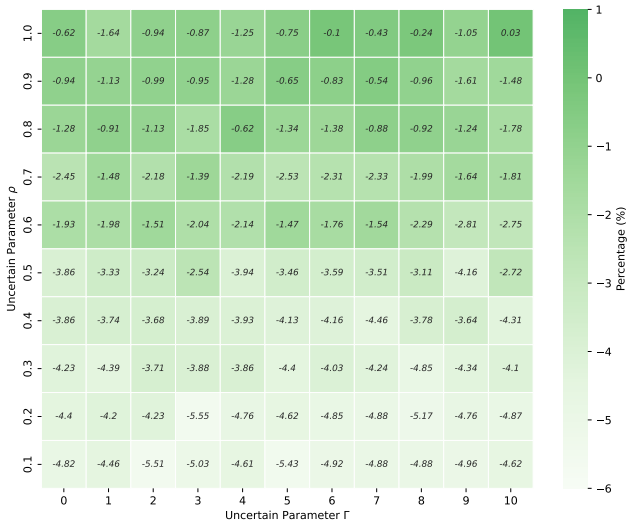

Fig. 8: Relative percentage reduction of average customer wait time for the robust MIVR model compared to predictive prescription with KNN-5.

Figure 10 displays the percentage decrease of average non-occupied VMT for the robust MIVR model compared to predictive prescription with KNN-5. When a certain level of demand uncertainty is considered in the robust MIVR model, it reduces the average non-occupied VMT for each vehicle.

Figure 11 exhibits the percentage reduction of average vehicle rebalancing trips for the robust MIVR model compared to predictive prescription with KNN-5. The robust MIVR model significantly reduces the number of rebalancing trips dispatched in the system. Given that robust optimization generates decisions optimal for the worst-case scenario, the robust MIVR model is conservative and few rebalancing trips are made to mitigate the impact of inaccurate demand estimations. On the other hand, predictive prescriptions generate decisions that are optimal for an expected scenario which indicates demand from

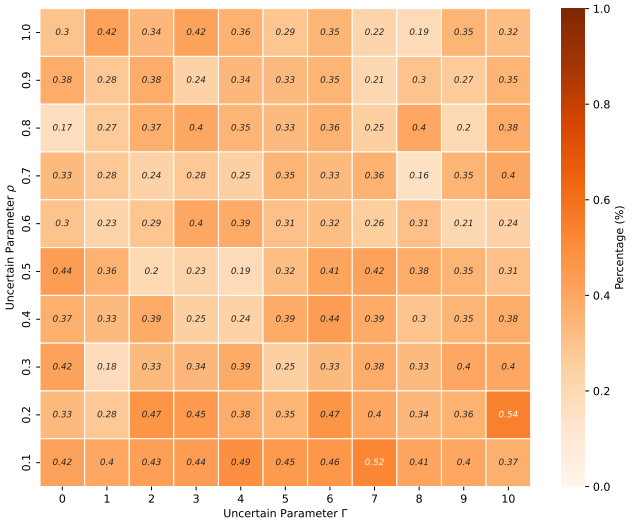

Fig. 9: Absolute reduction of customer unsatisfied rate for the robust MIVR model compared to predictive prescription with KNN-5.

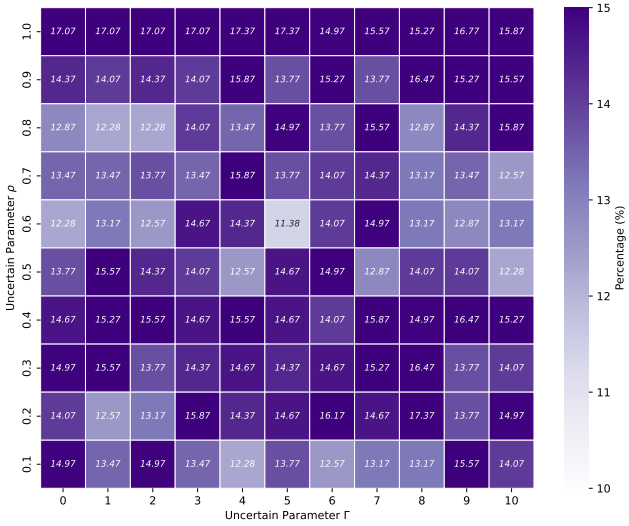

Fig. 11: Relative percentage reduction of average rebalancing trips for the robust MIVR model compared to predictive prescription with KNN-5.

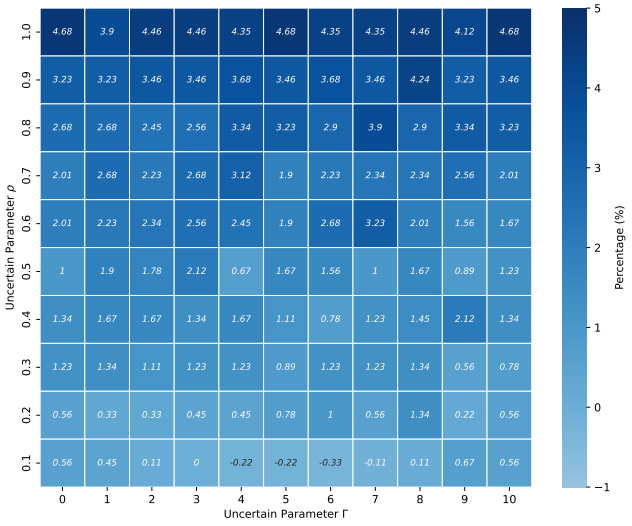

Fig. 10: Relative percentage decrease of average non-occupied VMT for the robust MIVR model compared to predictive prescription with KNN-5.

previous m days. Therefore, they do not maintain the same level of conservativeness as the robust MIVR model.

In conclusion, the robust MIVR model satisfies more customers while conducting fewer rebalancing trips and predictive prescriptions reduce the average customer wait time. From a practical perspective, applying the robust MIVR model requires decision-makers to choose an uncertainty level (ρ and Γ) incorporated in the model for the future demand. While for predictive prescriptions, additional information about the future demand is not required to make rebalancing decisions. Decision-makers should choose the appropriate model based on data availability and confidence about the level of uncertainty in the future demand.

VI. Conclusions

In this paper, we introduce a novel data-driven optimization approach, predictive prescriptions, into the vehicle rebalancing problem to handle demand uncertainty in the ride-hailing system. Building upon a state-of-the-art vehicle rebalancing model, MIVR proposed by Guo et al. [10], point-prediction-driven optimization models and data-driven optimization models are proposed to improve the model performance against demand uncertainty.

Regarding point-prediction-driven optimization models, a graph convolutional LSTM and two zone-based LSTM models are constructed in this paper to predict future demand for each sub-region. As for data-driven optimization models, SAA and predictive prescription with KNN and ORT are introduced in this paper. A real-world simulation with NYC data is used to evaluate performances for point-prediction-driven optimization models, data-driven optimization models, and two benchmark models, optimization with historical average and optimization with true demand, under four different demand scenarios.

Between the data-driven optimization and point-prediction-driven optimization models, one should make a decision based on supply to demand ratio and the prediction accuracy. When the future demand can be predicted accurately, point-prediction-driven optimization models should be adopted. When the demand is volatile and hard to predict, data-driven optimization models perform better. The system performances can be further improved for data-driven optimization models when the supply to demand ratio is higher, indicating more idle vehicles are available to be redistributed. Among all data-driven optimization methods, predictive prescriptions perform better by leveraging auxiliary information.

Meanwhile, prediction errors over the future demand in the vehicle rebalancing problem can be beneficial to system performances when errors come from demand underestimation. The "conservativeness" brought by the demand under-

estimation improves the system performance due to highly uncertain demand in the future. The strong performances of the robust MIVR model proposed in [10] are also brought by the "conservativeness" embedded in robust models. However, predictive models usually aim for "unbiasedness", and weights overestimation and underestimation equally. A possible future research direction is to develop predictive models for ride-hailing systems which have an asymmetric loss function that favors underestimation over overestimation. Meanwhile, extra benefits brought by conservativeness due to demand underestimation should have a limit. Future research could identify such underestimation level where the vehicle rebalancing benefits the most.

The best-performing data-driven optimization model, predictive prescription with KNN-5, is also compared with the robust MIVR proposed in [10], which utilizes robust optimization techniques to protect rebalancing decisions against demand uncertainty. The robust MIVR model reduces the customer unsatisfaction rate while conducting fewer vehicle rebalancing trips. On the other hand, predictive prescriptions reduce the average customer wait time but serve fewer customers. In practice, the robust MIVR model should be utilized if knowing the demand uncertainty level in the future. In general, predictive prescriptions can generate competitive rebalancing decisions without knowing any additional future demand information. Another future research direction can be introducing data-driven robust optimization techniques into the MIVR model, which combines the benefits of both data-driven optimization and robust optimization.

From a practical perspective, rebalancing models need to be selected ahead of schedule. When considering a whole day's demand, demand uncertainty and prediction accuracy of predictive models change from time to time. Therefore, a good operation strategy is to separate the whole operation period into high and low uncertainty periods based on historical demand data. For low uncertainty periods, point-prediction-driven optimization models should be adopted. As for high uncertainty periods, data-driven optimization models, including robust and predictive prescription models, can be applied.

APPENDIX

A. Driver-Customer Matching Problem

In this section, the driver-customer matching problem utilized in the simulation for evaluating the performances of vehicle rebalancing models is described. Given locations for available vehicles $\mathcal{V} = \{v_1, ..., v_m\}$ and locations for customers who have requested a demand $\mathcal{R} = \{r_1, ..., r_n\}$, a driver-customer matching problem is solved to assign customer requests to drivers. Between a customer r_i and a vehicle v_j, let $d(r_i, v_j)$ and $\tau(r_i, v_j)$ represent the distance and travel time for picking up the customer, respectively. A customer will leave the system if the customer is not assigned to any drivers within the maximum wait time $\tilde{w}$.

To solve the driver-customer matching problem, we first construct a bipartite graph $G = (V, E)$, where $V = \mathcal{R} \cup \mathcal{V}$ and $E = \{e(r_i, v_j) : \forall r_i \in \mathcal{R}, \forall v_j \in \mathcal{V}, \tau(r_i, v_j) \leq \bar{w}\}$, indicating that an edge exists between a vehicle and a customer if the customer can be picked up by the vehicle within the maximum pickup time. The cost of each edge $e(r_i, v_j)$ equals to the pickup distance, i.e., $c_{e(r_i,v_j)} = d(r_i, v_j)$. The decision variables for the driver-customer matching problem are $x_{e(r_i,v_j)} \in \{0, 1\}$ for each edge $e(r_i, v_j) \in E$ in the bipartite graph G, and $y_{r_i} \in \{0, 1\}$ for each customer $r_i \in \mathcal{R}$. $x_{e(r_i,v_j)} = 1$ represents that the customer r_i will be picked up by the vehicle v_j in the optimal matching. $y_{r_i} = 1$ denotes that the customer r_i can not be satisfied. Let $\mathcal{I}(r_i)$ represent the set of edges connected to a customer vertex r_i in G. Similarly, let $\mathcal{I}(v_j)$ indicate the set of edges connected to a driver vertex v_j in G. The optimal driver-customer matching problem is formulated as:

$$\min \quad \sum_{e(r_i,v_j)\in E} c_{e(r_i,v_j)} x_{e(r_i,v_j)} + \gamma \cdot \sum_{r_i \in \mathcal{R}} y_{r_i} \tag{13a}$$

$$\text{s.t.} \quad \sum_{e(r_i,v_j)\in \mathcal{I}(v_j)} x_{e(r_i,v_j)} \leq 1 \quad \forall v_j \in \mathcal{V} \tag{13b}$$

$$\sum_{e(r_i,v_j)\in \mathcal{I}(r_i)} x_{e(r_i,v_j)} + y_{r_i} = 1 \quad \forall r_i \in \mathcal{R} \tag{13c}$$

$$x_{e(r_i,v_j)} \in \{0, 1\} \quad \forall e(r_i, v_j) \in E \tag{13d}$$

$$y_{r_i} \in \{0, 1\} \quad \forall r_i \in \mathcal{R} \tag{13e}$$

The objective function (13a) minimizes a generalized cost for the driver-customer matching which consists of total pickup distance and penalties for unsatisfied customers. γ is the penalty parameter for each unsatisfied customer. Constraints (13b) guarantee that each vehicle can only be matched with at most one customer. Each customer is either assigned to a vehicle or remained to wait in the system, which is ensured by constraints (13c). Constraints (13d) and (13e) make sure that the decision variables are binary.

B. Uncertainty Set in the Robust MIVR Model

In this section, we briefly describe the uncertainty set utilized in the robust MIVR model, more details can be found in [10]. Given the future demand r_i^k of sub-region i at time k, the uncertainty set in the robust MIVR model consists of a box uncertainty set controlled by the parameter ρ and a polyhedral uncertainty set defined by the parameter Γ.

Let μ_i^k and σ_i^k represent the mean and standard deviation of the demand in sub-region i at time k from the historical data, respectively. The box uncertainty set $\tilde{\mathcal{U}}_i^k(\rho)$ is defined as

$$\tilde{\mathcal{U}}_i^k(\rho) = \left\{ r_i^k : \left| \frac{r_i^k - \mu_i^k}{\sigma_i^k} \right| \leq \rho \right\} \quad \forall i \in N, \ \forall k \in K,$$

where ρ indicates the parameter controlling the difference between historical average demand and future demand for each sub-region i at time k.

The polyhedral uncertainty set $\bar{\mathcal{U}}^k(\Gamma)$ is defined as

$$\bar{\mathcal{U}}^k(\Gamma) = \left\{ (r_1^k, ..., r_n^k) : \left| \sum_{i=1}^{n} (r_i^k - \mu_i^k) \right| \leq \Gamma \right\} \quad \forall k \in K,$$

where Γ denotes the parameter ensuring that the overall changes across all sub-regions at time k should not exceed Γ. Then the complete uncertainty set $\mathcal{U}$ used in the robust MIVR model is

$$\mathcal{U} = \left[\bigcap_{i=1}^{n}\bigcap_{k=1}^{\kappa}\tilde{\mathcal{U}}_i^k(\rho)\right] \cap \left[\bigcap_{k=1}^{\kappa}\mathcal{U}^k(\Gamma)\right].$$

Larger values of ρ and Γ lead to larger uncertainty set in the robust MIVR model, which leads to more conservative rebalancing decisions. Besides, decentralized vehicle rebalancing systems were proposed as contingency plans when AVs lost connections with central dispatch systems. Chen et al. [26] proposed a decentralized cooperative cruising method for offline operations of AMoD fleets. Their proposed method shows significant performance improvements compared to strategies with random-selected destinations for rebalancing AVs under different fleet sizes.

ACKNOWLEDGMENT

This work was supported in part by the MIT Institute for Data, Systems, and Society (IDSS) Seed Funds. The authors would like to thank Nicholas S. Caros for generating input data used in experiments.

REFERENCES

[1] S. Shaheen, A. Cohen, B. Yelchuru, and S. Sarkhili, "Mobility on demand operational concept report," no. FHWA-JPO-18-611, Sep 2017. [Online]. Available: https://rosap.ntl.bts.gov/view/dot/34258

[2] B. Schaller, "Unsustainable? The Growth of App-Based Ride Services and Traffic, Travel and the Future of New York City," 2017. [Online]. Available: http://www.schallerconsult.com/rideservices/unsustainable.pdf

[3] J. Jiang, "More Americans are using ride-hailing apps," January 2019. [Online]. Available: https://www.pewresearch.org/fact-tank/2019/01/04/more-americans-are-using-ride-hailing-apps/

[4] E. Brown, "The Ride-Hail Utopia That Got Stuck in Traffic," February 2020. [Online]. Available: https://www.wsj.com/articles/the-ride-hail-utopia-that-got-stuck-in-traffic-11581742802

[5] G. Zardini, N. Lanzetti, M. Pavone, and E. Frazzoli, "Analysis and control of autonomous mobility-on-demand systems: A review," *arXiv:2106.14827 [cs, eess]*, Jun 2021, arXiv: 2106.14827. [Online]. Available: http://arxiv.org/abs/2106.14827

[6] K. Spieser, S. Samaranayake, W. Gruel, and E. Frazzoli, "Shared-vehicle mobility-on-demand systems: a fleet operator's guide to rebalancing empty vehicles," in *Transportation Research Board 95th Annual Meeting*, no. 16-5987. Transportation Research Board, 2016.

[7] A. Wallar, M. Van Der Zee, J. Alonso-Mora, and D. Rus, "Vehicle rebalancing for mobility-on-demand systems with ride-sharing," in *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2018, pp. 4539–4546.

[8] F. Miao, S. Han, A. M. Hendawi, M. E. Khalefa, J. A. Stankovic, and G. J. Pappas, "Data-driven distributionally robust vehicle balancing using dynamic region partitions," in *Proceedings of the 8th International Conference on Cyber-Physical Systems - ICCPS '17*. ACM Press, 2017, p. 261–271. [Online]. Available: http://dl.acm.org/citation.cfm?doid=3055004.3055024

[9] J. Wen, J. Zhao, and P. Jaillet, "Rebalancing shared mobility-on-demand systems: A reinforcement learning approach," in *2017 IEEE 20th International Conference on Intelligent Transportation Systems (ITSC)*. IEEE, 2017, pp. 220–225.

[10] X. Guo, N. S. Caros, and J. Zhao, "Robust matching-integrated vehicle rebalancing in ride-hailing system with uncertain demand," *Transportation Research Part B: Methodological*, vol. 150, pp. 161–189, 2021.

[11] H. Yao, F. Wu, J. Ke, X. Tang, Y. Jia, S. Lu, P. Gong, Z. Li, J. Ye, and D. Chuxing, "Deep multi-view spatial-temporal network for taxi demand prediction," *32nd AAAI Conference on Artificial Intelligence, AAAI 2018*, pp. 2588–2595, 2018.

[12] J. Ke, H. Zheng, H. Yang, and X. M. Chen, "Short-term forecasting of passenger demand under on-demand ride services: A spatio-temporal deep learning approach," *Transportation Research Part C: Emerging Technologies*, vol. 85, no. June, pp. 591–608, 2017. [Online]. Available: http://dx.doi.org/10.1016/j.trc.2017.10.016

[13] J. Ke, S. Feng, Z. Zhu, H. Yang, and J. Ye, "Joint predictions of multi-modal ride-hailing demands: A deep multi-task multi-graph learning-based approach," *Transportation Research Part C: Emerging Technologies*, vol. 127, p. 103063, 2021. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0968090X21000905

[14] X. Geng, Y. Li, L. Wang, L. Zhang, Q. Yang, J. Ye, and Y. Liu, "Spatiotemporal Multi-Graph Convolution Network for Ride-Hailing Demand Forecasting," *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 33, pp. 3656–3663, 2019.

[15] J. Ye, J. Zhao, K. Ye, and C. Xu, "How to Build a Graph-Based Deep Learning Architecture in Traffic Domain: A Survey," *IEEE Transactions on Intelligent Transportation Systems*, 2020.

[16] J. R. Birge and F. Louveaux, *Introduction to Stochastic Programming*, ser. Springer Series in Operations Research and Financial Engineering. Springer New York, 2011. [Online]. Available: http://link.springer.com/10.1007/978-1-4614-0237-4

[17] D. Bertsimas and N. Kallus, "From predictive to prescriptive analytics," *Management Science*, vol. 66, no. 3, pp. 1025–1044, 2020.

[18] H. Wang and H. Yang, "Ridesourcing systems: A framework and review," *Transportation Research Part B: Methodological*, vol. 129, p. 122–155, Nov 2019.

[19] G. A. Godfrey and W. B. Powell, "An adaptive dynamic programming algorithm for dynamic fleet management, i: Single period travel times," *Transportation Science*, vol. 36, no. 1, p. 21–39, Feb 2002.

[20] Y. Jiao, X. Tang, Z. Qin, S. Li, F. Zhang, H. Zhu, and J. Ye, "Real-world ride-hailing vehicle repositioning using deep reinforcement learning," *arXiv:2103.04555 [cs]*, Jul 2021, arXiv: 2103.04555. [Online]. Available: http://arxiv.org/abs/2103.04555

[21] A. Braverman, J. G. Dai, X. Liu, and L. Ying, "Empty-car routing in ridesharing systems," *Operations Research*, vol. 67, no. 5, p. 1437–1452, Sep 2019.

[22] M. Pavone, S. L. Smith, E. Frazzoli, and D. Rus, "Robotic load balancing for mobility-on-demand systems," *The International Journal of Robotics Research*, vol. 31, no. 7, p. 839–854, Jun 2012.

[23] R. Zhang and M. Pavone, "Control of robotic mobility-on-demand systems: a queueing-theoretical perspective," *arXiv:1404.4391 [cs]*, Apr 2014, arXiv: 1404.4391. [Online]. Available: http://arxiv.org/abs/1404.4391

[24] R. Iglesias, F. Rossi, K. Wang, D. Hallac, J. Leskovec, and M. Pavone, "Data-driven model predictive control of autonomous mobility-on-demand systems," *arXiv:1709.07032 [cs, stat]*, Sep 2017, arXiv: 1709.07032. [Online]. Available: http://arxiv.org/abs/1709.07032

[25] M. Tsao, D. Milojevic, C. Ruch, M. Salazar, E. Frazzoli, and M. Pavone, "Model predictive control of ride-sharing autonomous mobility-on-demand systems," in *2019 International Conference on Robotics and Automation (ICRA)*. IEEE, May 2019, p. 6665–6671. [Online]. Available: https://ieeexplore.ieee.org/document/8794194/

[26] L. Chen, A. H. Valadkhani, and M. Ramezani, "Decentralised cooperative cruising of autonomous ride-sourcing fleets," *Transportation Research Part C: Emerging Technologies*, vol. 131, no. May, p. 103336, 2021. [Online]. Available: https://doi.org/10.1016/j.trc.2021.103336

[27] L. Al-Kanj, J. Nascimento, and W. B. Powell, "Approximate dynamic programming for planning a ride-hailing system using autonomous fleets of electric vehicles," *European Journal of Operational Research*, vol. 284, no. 3, pp. 1088–1106, 2020. [Online]. Available: https://doi.org/10.1016/j.ejor.2020.01.033

[28] M. Ramezani and M. Nourinejad, "Dynamic modeling and control of taxi services in large-scale urban networks: A macroscopic approach," *Transportation Research Part C: Emerging Technologies*, vol. 94, pp. 203–219, 2018. [Online]. Available: https://doi.org/10.1016/j.trc.2017.08.011

[29] A. J. Kleywegt, A. Shapiro, and T. Homem-de Mello, "The sample average approximation method for stochastic discrete optimization," *SIAM Journal on Optimization*, vol. 12, no. 2, p. 479–502, Jan 2002.

[30] A. Ben-Tal, L. El Ghaoui, and A. S. Nemirovskiĭ, *Robust optimization*, ser. Princeton series in applied mathematics. Princeton University Press, 2009.

[31] D. Bertsimas, V. Gupta, and N. Kallus, "Data-driven robust optimization," *arXiv:1401.0212 [math]*, Nov 2014, arXiv: 1401.0212. [Online]. Available: http://arxiv.org/abs/1401.0212

[32] X. Shi, Z. Chen, H. Wang, D.-Y. Yeung, W.-k. Wong, and W.-c. Woo, "Convolutional lstm network: A machine learning approach for

precipitation nowcasting," in *Proceedings of the 28th International Conference on Neural Information Processing Systems - Volume 1*, ser. NIPS'15. Cambridge, MA, USA: MIT Press, 2015, p. 802–810.
[33] T. N. Kipf and M. Welling, "Semi-supervised classification with graph convolutional networks," *5th International Conference on Learning Representations, ICLR 2017 - Conference Track Proceedings*, pp. 1–14, 2017.
[34] D. Bertsimas and J. Dunn, "Optimal classification trees," *Machine Learning*, vol. 106, no. 7, p. 1039–1082, Jul 2017.
[35] NYC Taxi and Limousine Commission. (2019) TLC Trip Record Data. "[Online; accessed 15-June-2020]". [Online]. Available: www.nyc.gov/html/tlc/html/about/trip_record_data.shtml
[36] Uber Technologies, Inc. (2019) Uber Movement, New York City travel speeds. "[Online; accessed 15-June-2020]". [Online]. Available: https://movement.uber.com/explore/new_york/speeds/
[37] J. Bezanson, A. Edelman, S. Karpinski, and V. B. Shah, "Julia: A fresh approach to numerical computing," *SIAM Review*, vol. 59, no. 1, pp. 65–98, 2017. [Online]. Available: https://epubs.siam.org/doi/10.1137/141000671
[38] I. Dunning, J. Huchette, and M. Lubin, "Jump: A modeling language for mathematical optimization," *SIAM Review*, vol. 59, no. 2, pp. 295–320, 2017.
[39] Gurobi Optimization, LLC, "Gurobi optimizer reference manual," 2020. [Online]. Available: http://www.gurobi.com

Xiaotong Guo received the bachelor's degree in traffic engineering from Tongji University and the master's degree in transportation systems engineering from Cornell University. He is currently pursuing the Ph.D. degree in Transportation with the Department of Civil and Environmental Engineering at MIT. His research interests include shared mobility operations, public transit network design, and transit-centric multimodal urban mobility system.

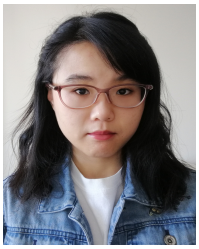

Qingyi Wang is a PhD student in transportation. She graduated with a Bachelor's of Applied Science in Engineering Science from the University of Toronto in 2018, and with a Master of Science in Transportation from MIT in 2020. Her research focuses on bringing machine learning into transportation demand modelling, as well as making use of domain knowledge to improve and better interpret machine learning models.

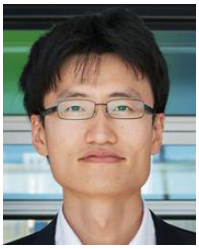

Jinhua Zhao is the Associate Professor of City and Transportation Planning at the Massachusetts Institute of Technology (MIT). He brings behavioral science and transportation technology together to shape travel behavior, design mobility system, and reform urban policies. He develops methods to sense, predict, nudge, and regulate travel behavior and designs multimodal mobility systems that integrate automated and shared mobility with public transport. He directs the JTL Urban Mobility Lab and Transit Lab at MIT and leads long-term research collaborations with major transportation authorities and operators worldwide, including London, Chicago, Hong Kong, and Singapore. He is the co-director of the Mobility Systems Center of the MIT Energy Initiative, and the director of the MIT Mobility Initiative.