A Data-Driven and Human-Centric EV Charging Recommendation System at City-Scale

Jingping Nie[1], Stephen Xia[1], Yanchen Liu[1], Shengxuan Ding[2], Lanxiang Hu[1], Minghui Zhao[1], Yuang Fan[3], Mohamed Abdel-Aty[2], Matthias Preindl[1], and Xiaofan Jiang[1]
[1]Columbia University, [2]University of Central Florida, [3]New York University
USA, {jn2551, stephen.xia, yl4189, lh3116, mz2866, matthias.preindl}@columbia.edu,
shengxuanding@knights.ucf.edu, yf1317@nyu.edu, M.Aty@ucf.edu, jiang@ee.columbia.edu

ABSTRACT

Electric vehicles (EVs) have gained widespread popularity in recent years, and the scheduling and routing of EV charging have impacted the welfare of both EV drivers and the grid. In this paper, we present a practical, data-driven, and human-centric EV charging recommendation system at the city-scale based on deep reinforcement learning (DRL). The system co-optimizes the welfare of both the EV drivers and the grid. We augmented and aggregated data from various sources, including public data, location-based data companies, and government authorities, with different formats and time granularities. The data includes EV charger information, grid capacity, EV driving behavior information, and city-scale mobility. We created a 30-day per-minute unified EV charger information dataset with charging prices and grid capacity, as well as an EV driving behavior dataset with location and State of Charge (SoC) information. Our evaluation of the recommendation system shows that it is able to provide recommendations that reduce the average driver-to-charger distance and minimize the number of times chargers switch to a different driver. The dataset we prepared for training the DRL agent, including augmented EV driving data and charging station information, will be open-sourced to benefit future research in the community.

CCS CONCEPTS

• **Social and professional topics** → *Sustainability*; • **Hardware** → *Smart grid*; • **Theory of computation** → *Reinforcement learning*.

KEYWORDS

Human-Centric, EV Charging, Recommendation System, Data-Driven, Smart Grid, Reinforcement Learning

ACM Reference Format:
Jingping Nie[1], Stephen Xia[1], Yanchen Liu[1], Shengxuan Ding[2], Lanxiang Hu[1], Minghui Zhao[1],, Yuang Fan[3], Mohamed Abdel-Aty[2], Matthias Preindl[1], and Xiaofan Jiang[1]. 2023. A Data-Driven and Human-Centric EV Charging Recommendation System at City-Scale. In *The 14th ACM International Conference on Future Energy Systems (e-Energy '23), June 20–23, 2023, Orlando, FL, USA.* ACM, New York, NY, USA, 12 pages. https://doi.org/10.1145/3575813.3597350

e-Energy '23, June 20–23, 2023, Orlando, FL, USA

ACM ISBN 979-8-4007-0032-3/23/06…$15.00
https://doi.org/10.1145/3575813.3597350

1 INTRODUCTION

The increasing popularity of EVs has led to a growing demand for charging infrastructure. With the increasing number of EVs on the road, the management of charging infrastructure becomes a critical issue for grid operators, charger operators, and EV owners. The charging infrastructure must be able to handle the increasing demand for energy, and the grid must be able to support the integration of EVs into the energy system. In this context, a data-driven city-scale human-centric EV-interfaced grid recommender system has the potential to provide efficient and sustainable charging solutions for EVs.

There are several sets of challenges that make this problem difficult. First, it is very difficult to simulate or model the driving path and speed of a real person. Additionally, abnormal fluctuations and errors in traffic flow occur due to unexpected traffic events [39]. Second, there is a physical constraint that needs to be considered when assigning EVs to chargers; Once an EV is assigned to the charger, there is a period of time, energy consumed, and distance traveled before the EV arrives at the charger. Finally, the total possible number of charger-EV assignments is huge. For example, in an environment with 160 chargers and 1,000 vehicles, there are $P(1000, 160) = \frac{1000!}{840!}$ possible charger-driver assignments.

A second set of challenges arises due to a lack of a unified dataset that we can use for training and modeling. There are a limited amount of datasets available for estimating vehicle SoC, EV charger characteristics, and power grid hosting capacity with fine time granularity. Additionally, publicly available EV mobility data is limited with low spatial and temporal resolution. Many of the vehicle mobility datasets only contain start and end points, rather than the actual path taken.

Previous works on EV charging management have focused on optimizing grid usage, maximizing the charging efficiency, or balancing the load between different charging stations. However, these works have limitations in terms of scalability, real-world applicability, and the ability to handle large amounts of data. In addition, prior studies often use power grid models, such as IEEE bus systems, rather than actual datasets from real-world scenarios. Moreover, they only focus on specific objectives such as reducing charging costs without taking into account the capacity of the grid.

In this paper, *we present a practical, data-driven, and human-centric EV charging recommendation system at the city scale based on deep reinforcement learning (DRL)* with three key features. First,

it builds on a carefully constructed city-scale EV driving dataset via the augmentation and aggregation of both public and self-collected datasets in different forms and with different granularity. Second, it takes into account practical considerations of a number of important parameters, including the varying EV demand, charging price, and grid capacity as a function of the status of the charging station at different timestamps. Third, instead of treating the drivers (EVs) as the agents in the DRL framework, which is considered in most of the recent works, we treat the chargers as the agents, as they are more directly related to important factors such as the grid capacity.

We create our recommender system, training environment, and simulation environment tailored to the Manhattan area because there exist a modest number of data sources that can be used to augment EV data from existing publicly available datasets. Additionally, the Manhattan area is one of several areas where data on power grid hosting capacity exists. Our proposed recommendation system leverages deep reinforcement learning (DRL) based on a deep Q-network (DQN) to make recommendations for optimized EV charging decisions, e.g., which EV each charger (agent) selects at a given time.

To build and evaluate practical recommendation systems, we construct a city-scale EV driving dataset via careful augmentation and aggregation of various datasets including diverse information such as vehicle trip routes, EV driving/charging sessions, and EV charger location and charging price. These data include both public datasets from different sources and a self-collected dataset with 3 volunteers using On-board diagnostics II (OBD-II) device driving a Tesla Model Y and Hyundai MPVs over the course of 3 months. Overall, our constructed EV driving dataset includes detailed labels ({timestamp, trip ID, location, speed, SoC}) of a total number of 61,392 (anonymized) valid trips over a duration of 30 days, and an EV charger dataset. The total dataset size is over 100 GB with second-level time granularity. Using this dataset, we train and test the DRL model, and demonstrate the effectiveness of the proposed EV charging recommendation system. For example, we show that the DRL-based recommendation system can, on average, reduce the distance between the EV chargers and their selected EV drivers by 32%. This dataset will be open-sourced and shared with the community. We believe the dataset we have generated and the recommender system framework we have proposed have the potential to act as an initial foundation for further investigations in the field of large-scale EV charging/routing optimizations.

To summarize, the main contributions of this paper include:

- We design an effective pipeline for pre-processing and augmentation of real-world EV driving data and charging infrastructure data, which closely simulate real-world charging processes.
- We create a realistic EV-driving-charging dataset with fine time granularity based on real-world data that, *for the first time*, combines the information from two domains: (i) the EV driving behaviors and (ii) the EV chargers. We open-source the dataset to facilitate the research in this area (https://github.com/Columbia-ICSL/Data-Driven-Human-Centric-EV-Charging).
- We propose and develop a practical, data-driven, and human-centric recommendation system for EV charging in Manhattan based on DRL, which takes into account the grid capacity, EV demand, and charging infrastructure when making recommendations. By considering each EV charger as an agent, the framework allows for each EV charger to make decisions about which set of EVs to serve when needed.
- We evaluate our proposed DRL-based recommendation system using the pre-processed and augmented dataset, demonstrating the effectiveness of the system in optimizing the charging of EVs.

In Section 4 and 5, we describe the preprocessing and augmentation of the real-world EV driving data and charging infrastructure data. In Section 6, we describe the DRL model used to make recommendations for the optimal charging of EVs. In Section 7, we evaluate the effectiveness of the system using real-world data. Finally, in Section 8, we discuss the limitations of the system and future work.

2 RELATED WORK

Our proposed framework, described in Section 3, takes input from 4 dimensions. In this section, we introduce related work in each aspect and summarize the corresponding existing datasets, traces, as well as simulation and optimization methods.

2.1 EV Charger & Charging Session Information

Typically, an EV drive relies on mobile apps and platforms such as *ChargeHub* [38], *chargepoint* [15], *NYserda* [6], *PlugShare* [7], and *TomTom* [14], to find available charging stations. Although these platforms provide users a list of charger locations with their pricing, charging rate, and availability information, they *do not* make any intelligent recommendation or provide a systematic way to corporately schedule users' charging behaviors. As a result, an EV driver usually selects the closest charging station or manually selects one from the list provided by the platforms.

For a network of EV chargers to cooperatively schedule users' charging behaviors, global information such as the EV charging sessions of each charger is required. Examples of publicly available EV charging session datasets include

- *Adaptive Charging Network Dataset (ACN-Data)*, which contains >32,000 EV charging sessions from charging stations near Caltech and NASA JPL, spanning multiple years [26];
- *Georgia Tech Electric Vehicles Charging Stations Dataset*, which contains >3,300 EV charging sessions from 105 campus-wide charging stations and 85 users, with second-level time granularity, taken over the course of one year [11].
- *Schneider Residential Charging Dataset*, which contains 6,878 EV charging sessions collected from residential chargers in apartment buildings in Norway, where loads are aggregated hourly [37].

Since none of these EV charging session datasets are collected from Manhattan, we crawl data from the above-mentioned mobile platform in Manhattan to be incorporated in our dataset (Section 5).

2.2 EV SoC Infomation

On the EV side, there are only a limited number of available datasets on the *Energy usage / SoC* information of EVs over time and vehicle location. The most comprehensive dataset is the *Vehicle Energy Dataset* [33], which contains timestamped data of fuel, energy, speed, and auxiliary power usage from 383 personal cars in the Ann

Arbor area, Michigan, including 264 gasoline vehicles, 92 hybrid electric vehicles, and 27 plug-in EVs. Estimating SoC is important for accurately simulating and informing the energy usage of EVs while driving and charging. Commonly used SoC estimation methods include the classical Coulomb counting method that calculates SoC from charging and discharging battery currents, the look-up tables mapping characteristic parameters of LIB to SoC [20], the open circuit voltage (OCV) method that seeks to precisely represent SoC as a non-linear function of OCV [10], the impedance method to fit SoC with ohmic resistance parameters [12], and the Ampere-Hour Integral method [40].

We combine the SoC change over time information collected from the 27 plug-in EVs in Vehicle Energy Dataset with the Coulomb counting-based SoC estimation to obtain the EV's SoC for each trip in our dataset (Sections 4.2–4.4). We select the Coulomb counting method since it is the most scalable due to a small number of parameters and easy access to battery currents in most open-source EV driving and charging datasets.

2.3 Vehicle Mobility

There are several data sources of vehicle driving and traveling patterns, including

- *National Household Travel Survey (NHTSA)* [1], which contains counts of people who traveled between different cities in the United States, the distance traveled, and the modality of transportation, from over 220 billion trips during 2017-2020;
- *NYC Taxi & Limousine Commission Trip Record Data* [8], which contains trip records from yellow and green taxis in the NYC area since 2009. This dataset contains pick-up/drop-off locations and times, fares, and distance traveled;
- *INRIX* is a company that provides real-time and historical data on traffic and road data from millions of GPS-enabled automobiles, conventional road sensors, and hundreds of other sources [22]. However, INRIX displays deviations and delays in speed and coordinates when vehicles are located in uncongested highways or congested urban areas[35];
- *HERE* is a company that also provides datasets for location-based and traffic data, in addition to a RESTful API for listing EV charging stations in a specific area [16];
- *Regional Integrated Transportation Information System (RITIS)* is a data-driven platform for transportation analysis, aggregating data from multiple agencies [19]. RITIS provides road information based on various factors, including road status, locations, and travel directions, collected from induction loops, side-firing sensors (acoustic, microwave, etc.), and radar.

In our dataset, we filter and augment information from these five data sources to generate vehicle mobility traces in Manhattan from 11/01/2022–11/30/2022 (Section 4.1).

2.4 Grid Capacity

With the fast development of Vehicle to Grid (V2G) technology and the growing presence of various distributed energy resources (DERs) such as solar photovoltaic units, battery storage systems, and EVs, researchers and grid operators are actively assessing and analyzing the hosting capacity of the power grid to limit the effects on power reliability and safety caused by adding more DERs [24, 30–32]. Although the grid is managed by around 1,600 grid operators, existing grid capacity maps cover only a few areas including NYC, New Jersey, Philadelphia, California, and Maryland [3]. These hosting capacity maps indicate how many generations (expressed in kW) can be added to a feeder before the feeder reaches capacity or other limitations that reduce the reliability of service. Works for pairing EVs with chargers for replenishing EV SoC using real-world data generally lack consideration for the power grid capacity.

To overcome the lack of grid information, researchers have built systematic simulators such as the *ACN-Simulator* [27], which is a tool designed to simulate the impact of EVs on the electricity grid and to test different charging management strategies. ACN-Simulator models the interactions between EVs, charging stations, and the electricity grid, allowing researchers to study the effects of different charging strategies on the overall energy demand and distribution. However, it only considers the scenarios after EVs arrive at the charging stations, but does not consider situations when the EVs are on the road and may not have charging plans.

We use the *EV Charging Capacity map* from Con Edison which contains relevant information in Manhattan (Section 5).

2.5 EV Charging Optimization

In the area of scheduling and optimizing EV charging, various methods have been proposed to balance the energy demand between the grid and EVs, including centralized scheduling and decentralized scheduling [9]. In the central approach, an EV aggregator communicates with EV users to minimize energy charging costs and arranges the charging pattern for each EV directly by collecting charging data from the EVs and the electrical grid. A centralized scheduling strategy considers historical and statistical data, such as the charging time, grid capacity, and charging rate, to minimize peak loads on the grid and satisfy driver expectations. In the decentralized approach, each EV user makes decisions on when to charge, based on its own optimization criteria and parameters. The EV aggregator only aggregates the bids from each EV.

Despite the richness of techniques that people adopt to optimize smart grids and charging systems, many of them, including methods like model-based reinforcement learning (RL) and model-free RL, are data-hungry and require a substantial amount of real-world data to search for optimal policies and ensure the convergence of various (policy-based or value-based) methods. How to collect a sufficient amount of data then emerges as a challenge since it's impractical to obtain a reasonably large dataset from any existing smart grid infrastructure to be tested on, which is also known as the cold start problem in recommender systems [23]. Previous studies have shown effectiveness in explicit solutions to this problem with approaches including active learning [18]. Over the recent years, it has witnessed an increasing number of works devoted to actively constructing EV driving and charging datasets [26, 33, 37] for researchers to tackle various optimization problems on charging scheduling algorithms, energy management, charging station installation, frequent regulation, voltage control, etc.

Although there are a variety of datasets for informing vehicle energy usage, vehicle mobility, and EV charger characteristics, they

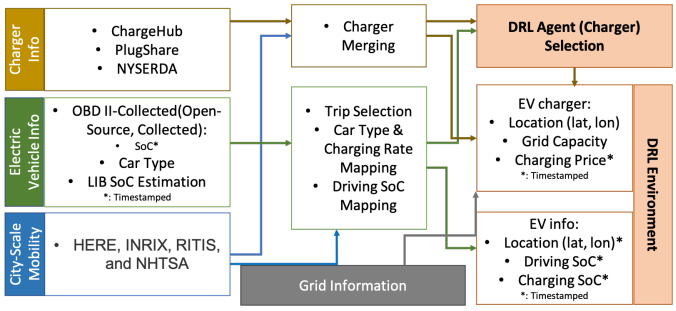

Figure 1: Data preparation and augmentation flow diagram of the proposed DQN-based DRL recommendation system.

are all very diverse. Specifically, there is no single dataset that integrates all information required to help model and train an intelligent recommender system. Moreover, existing works typically rely on power grid models, such as IEEE bus systems, instead of real-world datasets, and consider partial objectives such as simply minimizing the charging price without considering the grid capacity.

This work. Considering the aforementioned limitations and opportunities, we augmented and aggregated data from various sources, including public data, location-based data companies, and government authorities, with different formats and time granularity to create a *first-of-its-kind* 30-day per-minute unified *EV charger information dataset* with charging prices and grid capacity, as well as an *EV driving behavior dataset* with location and SoC information. To the best of our knowledge, this is the first work that considers the co-optimization of the welfare of both the EV drivers and the grid using a realistic dataset.

3 SYSTEM ARCHITECTURE

Due to the challenges mentioned in Section 1, we formulate our real-world, data-driven, human-centric recommender system for assigning EVs to chargers as a deep reinforcement learning (DRL) problem, where the *agent* is the collection of EV chargers. This is because the number of public EV chargers will be relatively constant compared to the number of drivers on the road throughout the day. More specifically, we create a deep Q-network (DQN) to realize this recommender system. There are two major components for the proposed recommender system: (i) data preparation and augmentation for the DRL environment, and (ii) training and applying the DQN agent. And the following assumptions are made:

(1) All EV drivers are rule-followers. Individual willingness, incentives, and unpredictable rule-breaking behaviors are not considered in this study.
(2) The system has a holistic view ("God's eye view") of the real-time states of all public-accessible EV chargers and EV drivers in the Manhattan area.
(3) The charging scheduling and EV assignment for private-owned EV chargers are not considered in this study.
(4) All the chargers in this study are Type-II chargers with a fixed charging rate of 10 kWh.

To obtain a *holistic* perspective, we consider four categories of data to model EV drivers and the effects of EV chargers on the grid, as shown in Figure 1 (more details in Section 4 and Section 5). The four categories include:

(1) Charger information (location, availability, and price) from ChargeHub, PlugShare, and NYSERDA;
(2) EV information, including state of charge (SoC), car type, charging rate, and discharging rate directly collected by On-board diagnostics II (OBD-II) devices from open-source datasets or collected in this study, the EV car type distribution [4], and LIB SOC estimation described in Section 4.3.1.
(3) City-scale vehicle mobility data includes the floating car data from HERE, INRIX, RITIS, and NHTSA.
(4) Grid information, which contains the network hosting capacity and EV charging capacity from Con Edison [2].

Because each EV charger data source provides slightly different information, we cleaned and merged data from these sources to obtain 520 EV charging stations with 1,061 total chargers to use in our system.

There are few publicly available open-source datasets containing *electric vehicle information*, namely SoC, and grid network hosting capacity. Because data for hosting capacity exists in the Manhattan, NYC area and because there exists a modest number of sources that can be used to generate vehicle mobility traces, **we decide to create our recommendation system for the Manhattan area**. In future work, we plan to explore and apply our system to other settings, as datasets from more regions become available.

To increase the amount of EV SoC data, we recruited three subjects driving EV cars (2 Tesla Model Y and 1 Hyundai MPV) in the Manhattan area using an onboard OBD-II device to collect driving and charging data for three months. Data from these three subjects are not representative enough of all EV driver behavior in Manhattan. As such, we leverage floating car data (speed, longitude (lon), latitude (lat)) to generate more mobility traces and vehicle trips, and we use the distribution of registered EV vehicles in Manhattan to assign a car type, charging rate, and discharging rate to these trips, as described in Section 4. After assigning the car type, we estimate the SoC of the vehicle during the trip using methods described in Section 4.3.

Considering the size of the state-action space, we select 160 chargers at 160 charging stations, as described in Section 5. Because many of the trips are extremely short in duration, we select only the 1,000 longest trips per day to train our model (Section 6), with 1 minute timesteps. The timestamped EV charger and EV driver SoC and mobility data are regarded as the environment for the DRL as depicted in Figure 1.

Although the EV chargers are regarded as the *agents*, the recommender system proposed in this work faces the EV drivers. Although there are open-source and microscopic road traffic simulation packages, such as SUMO [28], to simulate various traffic scenarios, including the movements and charging patterns of EVs. However, SUMO is majorly facing the developers, which is not suitable to run on a smartphone that EV drivers can use in their daily life. As such, we designed a smartphone system for EV drivers, shown in Figure 2. EV drivers input their origin, destination, vehicle type, current SoC, and willingness to be tracked at the beginning of each trip. If the user is not willing to be tracked, the route will be

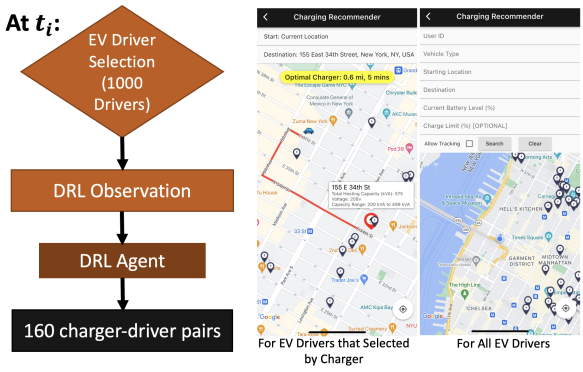

Figure 2: (Left) the process of making recommendations to EV drivers at each timestep, and (right) example screenshots of a customized smartphone app designed for the users using our proposed DRL-based recommendation system.

generated based on the Google Maps API, and SoC will be mapped by the method in Section 4.3.1.

At each timestep $t = t_i, i \in [0, 1, \cdots, 1,439]$, the 1,000 EVs that most likely need to be charged will be selected to form the *observation* or state that the EV charger agents will observe. The DRL agent learns a policy for choosing which of the 1,000 drivers should be charged at each of the 160 chargers at time t_i. Selected EV drivers for charging will receive a pop-up notification that will lead the driver to the corresponding charger.

4 VEHICLE DATA PREPARATION

In this section, we describe the real-world datasets, followed by the pipeline for the pre-processing and augmentation of these datasets for providing the *EV information* to be used by the DRL framework.

4.1 City-Scale Mobility Data and Augmentation

To obtain EV mobility traces in New York City, we aggregated datasets from three companies: INRIX, HERE, and RITIS (described in Section 2), which include the following traffic information: vehicle ID, traffic speed, ignition status, and vehicle coordinates.

We chose INRIX to gather the majority of traffic data, including vehicle coordinates and speed, since INRIX has recorded data from most number of vehicles, collected from hundreds of sources [22]. However, INRIX has many deviations, delays, and missing data in speed and coordinates in situations where the vehicles are located in uncongested highways or congested urban areas [35]. To supplement data defects from INRIX, we incorporated traffic speed and vehicle ignition data from HERE and RITIS.

Once the data from these various sources are combined into a unified format, we filtered the data according to [41] to ensure an appropriate quality of vehicle trajectories. To reduce noise and impute missing data in vehicle trajectories, we used moving averages and neighborhood averaging. To generate quantitative results, all study data were processed in Python on a server with one master node and four worker nodes. Each machine is outfitted with an AMD Ryzen thread ripper pro 3955wx 16-core and 32 processors, 256 GB of memory, and a 5TB disk capacity. We extracted vehicle traces in Manhattan, NYC, from 11/01/2022 to 11/30/2022.

To increase the number of vehicle traces, we leverage coarse-grained trip data from the NHTSA dataset [1], described in Section 2. This dataset only specifies origin and destination and does not specify vehicle speeds or trajectories. To generate trajectories, we adopt a similar method proposed in [21], by sampling road segments from origin to destination. Each road segment is weighted by how often they are utilized by vehicles in the dataset we collected from INRIX, HERE, and RITIS. *Overall, we construct a dataset with labels* ***{timestamp, trip ID, location (lat, lon), speed}*** *for a total number of* 61, 392 *anonymous valid trips (trips longer than 1 minute and move more than 0.5 m).*

4.2 Driving SoC Data

We used data from the Vehicle Energy Dataset (VED) [33] to model the energy consumption of EVs, and to extract the vehicle ID, speed, SoC, and distance traveled. Since the mobility data in VED is not from the NYC area, we did not incorporate this data into the mobility dataset we collected (Section 4.1). In this work, we assume that all EVs are non-hybrid and leave hybrid EVs for future work. Since the number of non-hybrid electric vehicles and trips is rather limited (only 27 vehicles with 2,888 trips), we collected data from three EV drivers (2 Tesla Model Y's and 1 Hyundai MPV) through onboard OBD-II loggers, over the course of three months (with 334 trips).

4.3 SoC Data Augmentation

As mentioned in Section 4.1, we augmented vehicle mobility traces using a publicly available travel dataset from the NHTSA. These newly generated mobility traces do not have an associated energy usage or SoC. Between the VED dataset and the data we collected from three additional EVs (see Section 4.2), we only have energy consumption and the SoC of 27 + 3 vehicles. Below, we describe how we generate new SoC traces for augmented vehicle mobility traces (Section 4.1), by first modeling the SoC of a vehicle's energy storage system (ESS) and then how a vehicle's mobility pattern affects power draw from the ESS.

4.3.1 LIB SoC Modeling. The battery models for electrical vehicles (BEV) play a central role in the model-based SoC estimation. Lithium-ion batteries (LIB) are the most commonly used ESS technology. SoC is a measure of the remaining energy in the LIB. To estimate SoC, the amount of remaining energy in the battery (ψ), which is defined as the average over maximum concentration of lithium-ions, is given by:

$$\psi = C_{s,avg}/C_{s,max}, \tag{1}$$

where $C_{s,avg}$ and $C_{s,max}$ represent the average and maximum concentration of lithium-ions in the LIB cathode, respectively.

Since is difficult to directly measure and model the lithium-ion concentration at any given moment to compute SoC, an alternative way to calculate SoC is by taking the ratio of current residual energy available to be drawn from the LIB, C_r, to highest possible load capacity, C_a, given by [36]

$$\text{SoC}(t) = C_r/C_a \cdot 100\%. \tag{2}$$

To associate the SoC with battery cell discharge current while driving, denoted by $I(t)$, we adopt the commonly used Coulomb Counting method, i.e. [17, 34]:

$$SoC\,(t) = \mathrm{SoC}(0) + \frac{\eta_i}{C_a} \cdot \int_0^t I(\tau)\ \mathrm{d}\tau, \tag{3}$$

where (*i*) $\mathrm{SoC}(0)$ denotes the initial state of charge, and (*ii*) η_i denotes the Coulombic Efficiency of the battery, or the ratio between the number of electrons extracted from the battery and the number of electrons put into the battery over a full cycle. This method allows us to easily derive the SoC at any timestamp with a configurable parameter C_a, which is based on the vehicle type, by measuring or simulating the current draw to and from the battery. Next, we discuss how we model the current draw from the LIB during a trip (Section 4.3.2) and while charging (Section 4.3.3).

4.3.2 Vehicle's SoC as a Function of Speed and Time. Next, we need to relate a vehicle's mobility (speed and distance traveled) to its SoC. We consider the following model of the vehicle's SoC at time t, $\mathrm{SoC}(t)$, as a function of its battery cell discharge current[1], $I(t)$, and the moving speed, $v(t)$,

$$\begin{aligned}\mathrm{SoC}(t) &= \mathrm{SoC}(0) + \frac{\eta_i}{C_a} \cdot \int_0^t I\,(\tau)\ \mathrm{d}\tau \\ &\approx \mathrm{SoC}(0) + \frac{\eta_i}{C_a} \cdot \int_0^t (a \cdot v(\tau) + b)\ \mathrm{d}\tau,\end{aligned} \tag{4}$$

where $I(t) = a \cdot v(t) + b$ is the first-order approximation of $I(t)$ as a function of $v(t)$. $P(t) = V_p \cdot I(t)$ relates the power consumption of the EV, $P(t)$, to the current draw and voltage of the battery pack, V_p. For LIB-powered EVs, we make the assumption that when speed is at cruising or lower speed with approximately $v < 100$ km/h, power consumption can be accurately approximated to the first order as shown in [25]. We only consider the case where the EV's LIB is never depleted to a level such that V_p degrades, i.e., $V_p \approx$ const. Differentiating both sides of (4) yields

$$\frac{\mathrm{dSoC}(t)}{\mathrm{d}t} \approx \frac{\eta_i}{C_a} \cdot a \cdot v(t) = r_v \cdot v(t), \tag{5}$$

where, $r_v := \frac{\eta_i a}{C_a}$ is the discharging rate in units of SoC/sec. We assume that an EV's energy consumption at rest is negligible, which eliminates the coefficient b due to the resulting boundary condition: $\frac{\mathrm{dSoC}(t)}{\mathrm{d}t}|_{v=0} = 0$. Using (5), the SoC of a vehicle at time t is given by

$$\mathrm{SoC}(t) = \mathrm{SoC}(0) + \int_0^t \frac{\mathrm{dSoC}(\tau)}{\mathrm{d}\tau}\ \mathrm{d}\tau \approx \mathrm{SoC}(0) + \int_0^t r_v \cdot v(\tau)\ \mathrm{d}\tau. \tag{6}$$

To obtain the discharging rate, r_v, we fit a linear regression model using data across 50 trips from real data collected over the course of 3 months from 2 volunteers driving the Tesla Model Y (Section 4.1). We discuss how we generate the discharging (r_v) and charging (Q) rates for the rest of the vehicles, for which we did not manually collect data, in Section 4.3.4. Table 1 summarizes these rates for each of the vehicle models.

4.3.3 SoC Charging Rate. When an EV is being charged at a charging station that provides a constant charging current of I_c, its SoC over time is given by

$$\begin{aligned}\mathrm{SoC}(t) = \mathrm{SoC}(0) + \frac{\eta_i}{C_a}\int_0^t I_c(\tau)\ \mathrm{d}\tau &= \mathrm{SoC}(0) + \eta_i \cdot \frac{I_c}{C_a} \cdot t \\ &= \mathrm{SoC}(0) + \eta_i \cdot Q \cdot t,\end{aligned} \tag{7}$$

where $Q := I_c/C_a$ denotes the charging rate, which is the ratio of the charging current, I_c, to the battery capacity, C_a. Intuitively, it takes longer time to charge a battery with larger capacity. If the battery capacity, C_a, can be found for the vehicle type, we can directly substitute this value. In cases where the value of C_a is unknown, we estimate the capacity from the vehicle's maximum range, $R_{\max}$, by learning a coefficient, α, from past vehicle trips [29], i.e.,

$$C_a = \alpha \cdot R_{\max} \Rightarrow Q = \frac{I_c}{C_a} = \frac{I_c}{\alpha \cdot R_{\max}}. \tag{8}$$

We show the charging rates in Table 1 and discuss how we generalize to other vehicle types in Section 4.3.4.

4.3.4 Generalization to other EV Models. To generalize the SoC to different vehicle types, we compare the battery capacity, C_a, and maximum vehicle range, R_{max}, between different vehicle models. From Section 4.3.2, we see that the mapping coefficients, r_v, learned for the Tesla Model Y is a function of C_a (we can ignore the coulombic efficiency, η_i, and the number of battery cells without loss of generality). We use the following equation to estimate the mapping coefficient, $r_{v,M}$ for vehicle type M based on the coefficient of the Tesla Model Y, $r_{v,\mathrm{Tesla}}$:

$$r_{v,M} = \frac{C_{a,M}}{C_{a,\mathrm{Tesla}}} \cdot r_{v,\mathrm{Tesla}}, \tag{9}$$

where $C_{a,M}$ denotes the battery capacity of vehicle type M.

Let $R_{\mathrm{max},Tesla}$ denote the maximum range for the Tesla Model Y and $R_{\mathrm{max},M}$ denote the maximum range for vehicle type M. For estimating the charging rate of vehicle type M, Q_M, using the charging rate of the Tesla Model Y, Q_{Tesla}, we can relate vehicle type M to the Tesla Model Y using the *capacity ratio* given by $q_M = R_{\mathrm{max},M}/R_{\mathrm{max,Tesla}}$. Together with (7), the SoC of vehicle M at time t is given by

$$\begin{aligned}\mathrm{SoC}_M(t) &= \mathrm{SoC}_M(0) + \eta_i \cdot \frac{1}{q_M} \cdot \frac{I_c}{\alpha R_{\mathrm{max,Tesla}}} \cdot t \\ &= \mathrm{SoC}_M(0) + \eta_i \cdot \frac{Q_{\mathrm{Tesla}}}{q_M} \cdot t.\end{aligned} \tag{10}$$

The full list of charging and discharging rates of the vehicle types we used in this work is shown in Table 1. We selected the top 10 most common EV types in NYC based on the NYSERDA EV Registration Map [6] and New York State DMV Records [13], where the "others" type is the average of the 10 selected vehicle models.

4.4 EV Driving Dataset for the DQN-based DRL Recommender

Based on the distribution of the 11 car types listed in Table 1 and the distribution of the initial SoCs of the 3,222 trips from VED dataset and the EV driving data collected by us, a car type and a starting SoC are assigned to each of 61,392 trips. With the car type and the starting SoC, using the method mentioned in Section 4.3.3 and the parameters listed in Table 1, the SoC for each timestamp is interpolated to the dataset prepared in Section 4.1 for the 61,392 trips. *Overall, with SoC data augmentation, we construct an EV driving dataset with labels* ***{timestamp (minute), trip ID, location (lat, lon), speed, SoC}*** *for a total number of 61,392 anonymous valid trips for a duration of 30 days, with a total dataset size of 10.3 GB.* This dataset is then used for the training and testing of the proposed DRL framework, described below.

[1]Typically, the total battery discharge current, $I(t)$, consists of the sum of currents generated by individual battery cells [25].

Table 1: Summary of maximum vehicle range (R_{max}), battery capacity (C_a), estimated discharging rate (r_v), estimated capacity ratio (q), and estimated charging rate (Q) for the 10 most common electric vehicle types in NYC.

Veh Brand	Veh Range (km)	LIB Capacity (kWh)	Discharging Rate (SoC/sec)	Capacity Ratio (q)	Charging Rate (SoC/sec)
Tesla Model Y	438	50	-3.860e-4	1.0	2.155e-3
BMW MPV	391	70.7	-3.450e-4	1.414	1.524e-3
BMW LDV	484	84	4.270e-4	1.68	1.283e-3
Audi MPV	357	95	-3.153e-4	1.9	1.134e-3
Ford MPV	398	70	-3.514e-4	1.4	1.540e-3
Kia MPV	499	77	-4.400e-4	1.54	1.401e-3
Volvo MPV	359	75	-3.172e-4	1.5	1.437e-3
Hyundai MPV	488	77.4	-4.301e-4	1.548	1.392e-3
Nissan LDV	338	60	-2.981e-4	1.2	1.796e-3
Others	376	73.2	-3.313e-4	1.465	1.472e-3

Figure 3: Location map of the chargers in Manhattan, NYC, that are selected as the *agents* in the proposed DRL framework. Green markers denote the chargers that are, on average, furthest away from the 61,392 vehicle mobility traces from our augmented dataset and black markers denote the chargers that are, on average, closest to these traces.

5 EV CHARGER DATA PREPARATION

To obtain the comprehensive dataset of EV charger information in Manhattan area, the charging station locations and information from five different sources: Plugshare, NYSERDA, Flo, Chargehub, and Tesla Supercharger, are crawled. Comparing the merge results and considering the Tesla chargers are not adaptable to all types of EVs, in this study, we acquire the real-time data from Plugshare, NYSERDA, and Chargehub. The chargers are merged based on their geographical locations. In addition, the price information is provided in different formats from these sources (e.g., "$/ kWh", "$/$hour$", and "same to the parking fee"). And some chargers miss the price information.

To simplify the environment for this recommender system, only Type-II chargers are filtered out and an assumption that all chargers charge the EV at 10 kW is made. With this charging rate and taking the median EV battery size 73.2kWh, all the prices with "$/$hour$" are converted into "$/ kWh". In addition, if the charging station misses the price information or is described in other formats, the real-time data from NYISO is used to generate the charging price in "$/ kWh" [5]. *Overall, we generate the timestamped charger information with the unified format ({**timestamp (minute), charger ID, location (lat, lon), charging price**}) for a total number of 1,061 chargers from 520 charging stations.*

As mentioned in Section 4.1, there are 61,392 trips in total over 30 days, or approximately 2,046 trips per day. In order to simulate the conditions that there will be more EV drivers that are actively driving on the road and has the potential needs of charging their cars, we select a total number of 160 chargers, where 80 chargers are on average closest to all of the vehicle trips (black markers in Figure 3), and the remaining 80 chargers are the furthest away from all of the vehicle trips (green markers in Figure 3).

Con Edison offers *Network Hosting Capacity* and *EV Charging Capacity* maps that are EV charging-related. The *Network Hosting Capacity* map provides the regional network area hosting capacity (kVA) in 6 levels. However, the hosting capacity around the 160 selected chargers in this study is almost the same. The *EV Charging Capacity* map provided by Con Edison provides the information for: (i) 460v Transformers, (ii) 208v Transformers, and (iii) Load Capacity for 3PH Feeders. As there is no Load Capacity information in Manhattan, it is assumed that the chargers with more 208v and 460v transformers around (within 1 km) brought fewer dynamics to the grid. The 30-day *EV Charger* information for the DRL environment is generated with **{timestamp (minute), trip ID, location (lat, lon), number of transformers around, charging price}**.

6 DEEP REINFORCEMENT LEARNING ARCHITECUTURE

Markov decision process (MDP) is usually used as the mathematical framework to describe an environment in DRL problems. Our human-centric real-world EV charging recommender system can be modeled as an MDP with discrete timestep (minutes in a day), t_i. We consider an EV charging recommendation system with 1,000 drivers, labeled by $d \in \mathcal{D} = \{1, 2, \ldots, 1,000\}$, and 160 charging stations (chargers), labeled by $c \in C = \{1, 2, \ldots, 160\}$. For each charger c at time t_i, it can select a driver $d \in C$ to provide charging service, or stay idle. We formulate a discounted MDP (with $\gamma = 0.99$) consisting of a 4-tuple: $\langle \mathbf{S}, \mathbf{A}, P_{\mathbf{a}}(\mathbf{s}(t_i), \mathbf{s}(t_{i+1})), R_{\mathbf{a}}(\mathbf{s}(t_i), \mathbf{s}(t_{i+1}))\rangle$:

- **S** represents the set of states, or the *space state*. In our case, the system state at t_i, denoted by $\mathbf{s}(t_i)$, consists of the location of

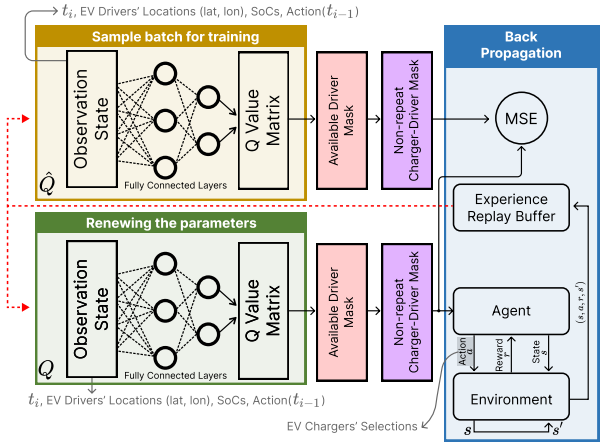

Figure 4: The proposed Deep-Q-Network- (DQN-) based deep reinforcement learning (DRL) architecture.

driver d at time t_i, $\text{lat}_d(t_i)$ and $\text{lon}_d(t_i)$ (latitude and longitude), the SoC of driver d, $\text{SoC}_d(t_i)$, the state of charging station c at t_{i-1}, and the current time t_i. As a result, the dimension of $\mathbf{s}(t_i)$ is $\mathbb{R}^{3,161\times 1}$.

- $\mathbf{A}$ represents the set of actions, or the *action space*. In particular, the action at time t_i, denoted by $\mathbf{a}(t_i) = [a_c(t_i)] \in \mathbb{N}^{160\times 1}$, where $a_c(t_i) \in \{\text{idle} \cup \mathcal{D}\}$ represents whether charging station c selects a driver or remains idle. We assume that a charger can select *at most* one driver at a time and any driver can be selected by *at most* one charger. Even with this assumption, since are a total number of 160 chargers and 1,000 EV drivers, the entire action space would be more than 160-permutations of 1,000 (the number of idling chargers can be more than one), which is more than 1.43×10^{474}. It is unrealistic to train a DQN agent with such a large action space. To address this issue, a *mask* is implemented after the Q value matrix is generated by the DQN to reduce the action space to only $160\times 1,001$, as shown in Figure 4. Specifically, we repeatedly choose the largest Q value in the Q value matrix, pair its corresponding charger c and driver d, and add a mask to this charger and driver, until the largest Q value corresponds to a charger, $c = m$, that chooses to be *idle*. Afterward, the charger, $c = m$, and the rest of the chargers that haven't been paired with a driver will be set to *idle*.
- $P_{\mathbf{a}}(\mathbf{s}(t_i), \mathbf{s}(t_{i+1}))$ denotes the transition probability that action $\mathbf{a}$ is taken in state $\mathbf{s}(t_i)$ will lead to state $\mathbf{s}(t_{i+1})$. This transition probability is complex, as it is highly correlated to the randomness of EV drivers on the road and their corresponding locations and SoCs. We use a model-free DQN-based approach to learn the transition procedure as illustrated in Figure 4. In particular, the DQN contains three fully connected MLP layers, with an input dimension of $3,161 \times 1$ and an output dimension of $160 \times 1,001$.
- $R_{\mathbf{a}}(\mathbf{s}(t_i), \mathbf{s}(t_{i+1}))$ denotes the reward received after transitioning from $\mathbf{s}(t_i)$ to $\mathbf{s}(t_{i+1})$ after action $\mathbf{a}(t_i)$ is taken. To ensure that the recommender system satisfies common preferences for EV drivers and learns policies that follow physical constraints when driving and charging, we impose several constraints on $R_{\mathbf{a}}$, as further discussed in Section 6.1.

In addition, as indicated in Figure 4, an *Available Driver Mask* is also applied to make sure that the chargers choose the drivers that are active (on the road). However, as the EV driver data is augmented with the floating car data, there is only trip information but no charging behavior included in the current environment. Additionally, many EV drivers are active for less than 1 hour, whereas general EV drivers usually charge their cars for more than 1 hour. Considering these issues, a *Driving Behaviour Information Remapping* strategy (see Section 6.2) is designed and applied on the drivers selected by the chargers as shown in Figure 1.

We use a dataset consisting of timestamped data collected from 1,000 drivers over 30 days with a 1-minute timing granularity. We select the data from the first 20 days to be used as the training set and the data from the remaining last 10 days to be used as the test set. The model is trained for 2,000 episodes, where each training episode spans the entire 20-day training set with $1,440 \times 20$ steps (1,440 minutes/day $\times$ 20 days), and it ends when the single-step reward is greater than 0. We set a learning rate of 3×10^{-4} and apply ϵ-greedy exploration with exploration probability $\epsilon = 0.2$.

6.1 Practical Modeling of the Reward Function

A unique feature of our proposed RL architecture is its ability to closely resemble real-world EV charging systems via carefully constructed constraints in the reward function. First, EV drivers won't detour a long way to charge their vehicles. Instead, they would usually charge along the way or at the source/destination of the trip. As such, the distance between each charger (*agent*) and its selected EV is analyzed. Consider 1-day time window, if $N_c(t_i)$ chargers select to provide charging service to $N_c(t_i)$ (distinct) drivers under action $\mathbf{a}(t_i)$, we set $R_{\text{distance}}(t_i) = -50 \cdot D(t_i)$, where $D(t_i)$ is the sum distance of all N_c charger-driver pairs (in km). Since the number of drivers on the road can exceed the number of chargers, the proposed DRL incorporate a reward, $R_{SoC}(t_i)$, to prioritize the drivers with lower SoC levels. Specifically, we consider a piecewise step function that respectively maps the SoC of each vehicle's, $\text{SoC}_d(t_i) \in [0, 50)/[50, 80)/[80, 100)/[100, +\infty)$, to a reward function, $R_{\text{SoC},d}(t_i) = 200/160/60/-1,000$. Then, when N_c charger-driver are paired, the total reward based on the individual vehicle's SoC is given by $R_{\text{SoC}}(t_i) = \sum_{d=1}^{N_c(t_i)} R_{\text{SoC},d}(t_i)$. To prevent chargers from always staying in the idle state, a negative reward of $R_{\text{idle}}(t_i) = -50 \cdot [160 - N_c(t_i)]$ is imposed.

In addition, since our proposed recommender system aims to co-optimize the welfare of the EV drivers with the grid capacity, it is encouraged that the drivers charge their vehicles chargers that are connected to a more stable grid location. In this work, we consider 208 V and 460 V transformers to use their voltage numbers to indicate the grid capacity. In particular, if $N_c(t_i)$ chargers are providing charging service, we set $R_{\text{grid}}(t_i) = 2 \cdot N_{T208}(t_i) + 5 \cdot N_{T460}(t_i)$, where $N_{T208}(t_i)$ and $N_{T460}(t_i)$ indicate the total number of 208 V and 460 V transformers around the $N_c(t_i)$ chargers.

Lastly, from the EV drivers' perspective, they would like to charge their vehicles at chargers with lower prices. If $N_c(t_i)$ chargers are paired with drivers at time t_i, we set a reward $R_{\text{price}}(t_i) = -100 \cdot \sum_{c=1}^{N_c(t_i)} Price_c(t_i)$, where $Price_c(t_i)$ indicates the charging price of charger c at time t_i. In addition, considering the fact that EV drivers are less likely to use their vehicles during night time,

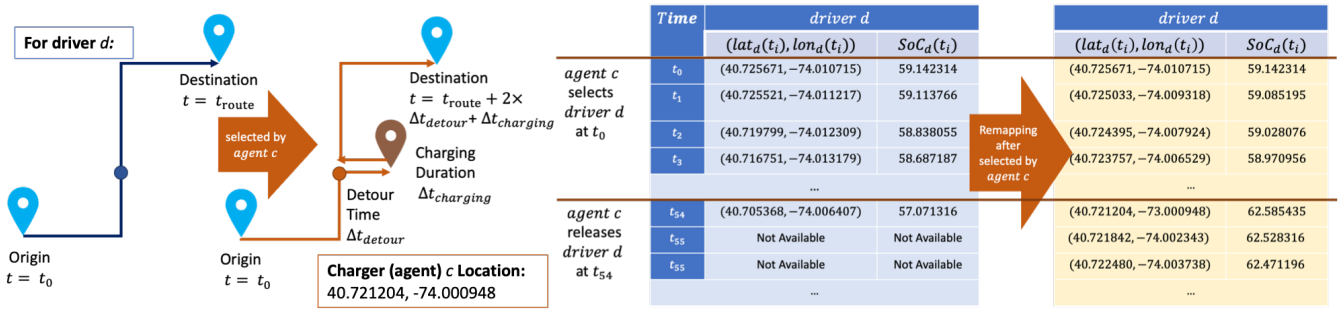

Figure 5: A demonstration of the driver location (lat, lon) and SoC remapping after selected by an *agent* (charger). In this example, the EV driver d will drive towards the charger (*agent*) c after it is selected, and will return back to its original route and continue the trip after it is released by charger c with $\Delta t_{detour} = 7$ min and $\Delta t_{charging} = 40$ min.

we also incorporate a piecewise switching cost penalty term as a function of the time of day. Specifically, we set $R_{\text{switching}}(t_i) = -Discount(t_i) \times N_c(t_i)$, where $Discount(t_i) = 1$, 150, and 200, for $t_i \in [360, 1380)$ (6am–11pm), $[1380, 1440) \cup [0, 59]$ (11pm–1am), and $[60, 360)$ (1am–6am), respectively.

Putting everything together, the total reward at t_i, $R_{\text{a}}(t_i)$, which is a function of the system state, $\mathbf{s}(t_i)$, and action, $\mathbf{a}(t_i)$, is given by

$$R_{\text{a}}(t_i) = R_{\text{distance}}(t_i) + R_{\text{SoC}}(t_i) + R_{\text{idle}}(t_i) + R_{\text{grid}}(t_i) + R_{\text{price}}(t_i) + R_{\text{switching}}(t_i) + \text{const.}, \quad (11)$$

where const. is a constant that we empirically set to be 4.1×10^7.

6.2 Driving Behaviour Information Remapping

As shown in Figure 5, when an agent (charger) takes action and selects the vehicle, it will grab the vehicle at a particular location and timestamp, detouring it to the charging station. For example, when a vehicle begins its trip (*route*) from the *Origin* and the charger selects it at time t_{detour}, the car will be simulated to drive towards the charging station. The time it takes to reach the charging station is denoted as Δt_{detour}. It may stay at the charging station until the agent performs a releasing action (selecting another driver or changing its state to idle), taking time $\Delta t_{charging}$. Then, the vehicle will follow the route of the detour back to where it left the original trip, taking another time Δt_{detour}, and continue the trip to the originally planned destination. In this way, the original locations for the rest of the trip to the *Destination* will be associated with an updated timestamp, i.e., $t_{route} \leftarrow t_{route} + 2 \cdot \Delta t_{detour} + \Delta t_{charging}$.

We interpolate the duration of this detour and the time at the charging station into the original route data. The timestamped driver location (lat, lon) and SoC of the rest of the original trip will be adjusted during the duration of the detour and charging and the environment in the DRL will be updated. Specifically, for each selected driver, based on its car type, current state (lon, lat, SoC), and distance to the corresponding charger, the discharging rate during driving (r_v) and the charging rate (Q) listed in Table 1 are used to interpolate the change in SoC of this driver's EV while driving to the charger, charging at the charging station, and getting back to the original trace. Google Map API is used to generate the route and moving speed. Figure 5 shows *an example* of how the location and SoC of driver d's vehicle change if charger c selects it at t_0 and releases it at t_{54}, with $\Delta t_{detour} = 7$ min and $\Delta t_{charging} = 40$ min. In this way, we have tuned the driver's behavior based on floating car data with reasonable charging behaviors (e.g., driving to the charger, charging at the station, and getting back to the original destination). These changes for every driver selected by each charger are constantly updated in the DRL environment.

Figure 6: DRL training reward across the episodes.

7 EVALUATION

In this section, we evaluate the effectiveness of the DRL model trained using the environment containing the *EV Charger* and *EV driver* information described in Section 4 and 5. We first implement a baseline scheme and show its defects. Then, we evaluate the change in training reward over the number of training episodes, and we run the model on test data and show the results.

Baseline Scheme. We consider and implement a baseline scheme where each user selects the charging station closest to the destination of its trip to charge to full if the EV's SoC drops below 90%. This baseline mimics typical users' daily EV charging behaviors, e.g., a user charges the EV at the working place (after driving to work) or at home (after driving back home from work). We assume there is no switching cost between EVs at each charger. We show that this scheme can cause a backlog of EVs that need charging services for some busy charging stations, while our proposed DRL-based optimization scheme can completely avoid any backlogged

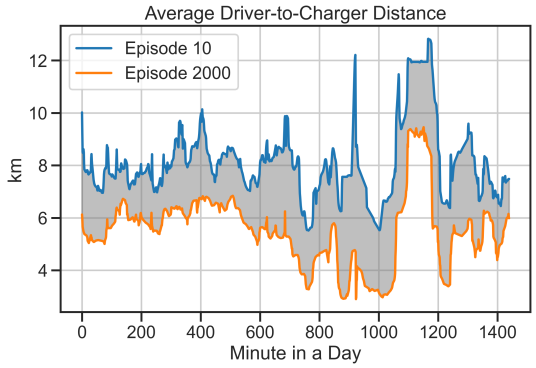

Figure 7: The average driver-to-charger distance (km) in one day for all driver-charger pairs generated by the DRL model trained after 10 (blue) vs. 2,000 (orange) episodes. The shadowed area indicates the improvements that the model made from episodes 10 to 2,000.

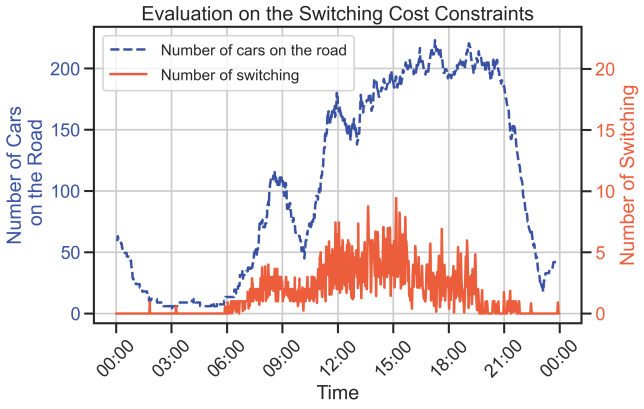

Figure 8: The number of drivers on the road during the day and the number of times that the 160 chargers switch to select a different EV driver.

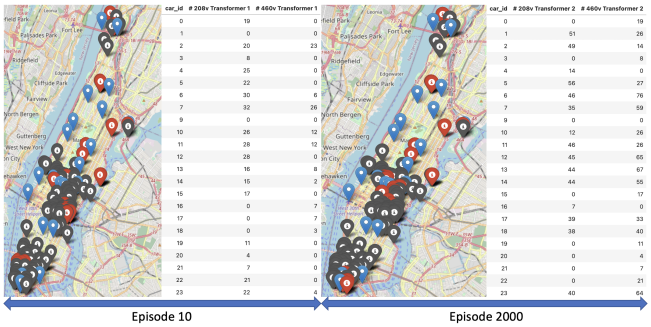

Figure 9: The difference in grid capacity information caused by different charger selections using the model saved in episode 10 (left) and episode 2,000 (right) at an example timestamp. Blue markers denote car location at that timestamp, black markers are all EV chargers and red markers denote chargers selected by the agent.

charger by using the *mask* introduced in Section 6. This is because by having the EV chargers as the agents, a charger will only "invite" drivers to charge when it is idle. As an example, we randomly select 1,000 EV drivers on 11/05/2022. Following the aforementioned rules, 611 users need to charge their EVs from the select 89 chargers (out of the 160 chargers). Under the baseline scheme, 44 chargers are selected by multiple EV drivers and a total number of 40 EV drivers need to wait for different periods of time before their EVs can be charged. In addition, since the baseline scheme does not consider the SoC of each EV, utilization of the EV chargers, grid stability, charging price, or costs associated with switching between EVs at each EV charger, it achieves an average reward (defined in Eq. (6.1)) that is more than twice lower than that achieved by the proposed DRL-based approach using the same 10-day test data.

Training reward. Figure 6 shows the training reward of each episode during the training process. The training reward increases over the number of episodes and stabilizes after around 200 episodes. The fluctuations in the rewards are mainly a result of using epsilon-greedy action sampling. We set the jumpout condition for each episode to be when the reward is larger than 0, or at the natural end of the episode (1440 minutes/day × 10 days). The epsilon-greedy action sampling allows the agent to continue exploring different actions even after it has converged to a good policy.

Average driver-to-charger distance. Next, we use one-day data from the testing set and show the effectiveness of the model by comparing the results from the model after 10 episodes vs the model after 2000 episodes. Figure 7 compares the average driver-to-charger distance generated by the model trained after 10 episodes, with which trained after 2000 episodes. As we can see, on average, after 10 episodes, the driver-to-charger distance is 8.11km whereas, after 2,000 iterations, the distance is reduced to 5.48km, yielding a 32% improvements and shows that the model is learning to make more efficient recommendations over the episodes.

Switching cost. The number of times that the EV chargers switch to select a different driver is another important metric to measure the effectiveness of the model. Figure 8 shows the number of drivers on the road in the test set during a day in the blue line, and the number of times the EV charger switches to a different user in the orange line. We can see that our agent has learned to efficiently find and charge EVs with low SoC when there are many cars on the road (during the daytime), resulting in a high number of EV chargers switching to select a different EV driver. Additionally, we can see that the time-related penalties we have implemented result in the charger-driver pairs not switching frequently during nighttime (11 PM–6 AM).

Grid capacity. Figure 9 shows the recommender's output using the model saved in episodes 10 and 2,000 using the same state information. We can see that as compared with episode 10, in episode 2,000, the model suggests chargers that have more substations around (transformers at 208V and 460V), as we discussed in Section 5.

8 LIMITATION, DISCUSSION AND FUTURE WORK

In this work we aggregated, cleaned, and augmented EV driver and EV charger data from various data sources (open-sourced,

location-based data companies, government authorities, and locally-collected), as well as proposed a DQN-based DRL recommendation system to provide charging recommendations for the EV drivers. Below, we summarize some of the current limitations and propose plans for future improvements and visions.

More Real EV Driver and Charger Data: Currently, we only have the 3-month EV driving data in New York City collected by 3 volunteers using OBD-II device. To generate all EV driving data used in this study, we augmented and simulated EV drivers in Manhattan using various real-world floating car data and EV driving data from the open-source dataset (VED) [33] and 3 volunteers, as well as various mapping techniques to assign the SoC and vehicle brands. We assume all EV drivers in this work are rule-followers.

However, if we could incorporate more data from a diverse group of real EV users, we can infer the travel habits and collect the preference of the specific user to provide more personalized recommendations. For example, if we could predict the user's current trip type based on their historical behavior, such as home-to-work/work-to-other, we would know that the driver might park the EV for a relatively longer time at home or at work than at a restaurant. In this case, we can provide recommendations that the EV drivers will find more reasonable and are more likely to accept. In addition, our dataset does not include situations where some EV users might have more than 1 EV. With more data from EV users, we could have the recommender take this into consideration.

We currently assume all the publicly available chargers as Type II chargers with a fixed charging rate of about 10 kW. We merged and augmented charging price and charger availability data from ChargeHub, NYserda, and PlugShare. And we filled the missing charging price with real-time electricity prices. Type II chargers have different charging rates among different brands and are installed at different locations. In addition, real EV chargers usually have varying charging rates during each charging session due to fluctuations in the grid. We are in the process of communicating with charger operators to acquire access to comprehensive real charger data (including timestamped charging rate) in addition to the connection time, charging time, etc., that are provided by the ACN dataset.

With more EV driver, charger accessibility, and price data, we can perform small-scale real-world evaluations to test out the performance of the proposed recommendation system.

More Real-time Grid Capacity Data: Currently the grid capacity information is acquired from publicly accessible maps as described in Section 5, which have low spatial and temporal granularity. We plan to get in touch with the grid operator for more grid capacity information with higher spatial-temporal granularity. As such, more precise hosting/EV charging capacity constraints can be added to the proposed framework and investigate how the charging behavior of EV drivers affects the resilience of the grid in the city scale.

Data Security, Privacy and Personalization: The collection and usage of personal EV driving data with GPS locations, EV statistics, and timestamps raise serious concerns about data security and privacy. To ensure the privacy and security of EV drivers' information, the data collection process and the proposed recommender system should be designed with privacy-by-design principles in mind.

One approach to address this concern is to implement a decentralized model, where the EV drivers' data is only stored locally on the edge device (e.g., the EV or EV driver's smartphone), rather than being transmitted to a central server. These data include drivers' historical trips, driving habits, user preferences, etc. The central server should only receive a processed version of these data and the current trip's information. We intend to integrate federated learning techniques into the decentralized model to make the recommender system more personalized for each EV driver by considering their preferences and past habits.

9 CONCLUSION

In this paper, we proposed and created a new, realistic dataset for a city-scale EV driving and charging, which combines information from both EVs and EV chargers, as well as the underlying power grid. Note that previous datasets only included partial information and were not collected within the same area. In addition, we presented a data-driven and human-centric EV charging recommendation system at the city scale, which optimizes EV charging while balancing grid usage, considering the welfare and behaviors of EV drivers. The system was trained using a DRL approach based on a deep Q-network and was evaluated using a carefully constructed EV charging dataset. The results showed that the proposed system was able to effectively make recommendations. For example, it reduced the average driver-to-charger distance by 32% as well as the switching cost of the chargers.

Although the proposed dataset can be further improved with more real-world EV driver and charger data, as well as higher spatial and temporal granularity of grid capacity information, we envision that our generated dataset and proposed recommender system framework can serve as an important stepping stone toward future research in the area of large-scale EV charging, EV routing, and EV charger dispatching optimizations. We plan to implement small-scale real-world experiments with real EV drivers and EV chargers to validate and improve the proposed DRL-based optimization scheme. We also plan to investigate decentralized approaches to ensure the privacy and security of the system. To benefit the community, the constructed EV charger information and EV driving behavior datasets used in this work were open-sourced, providing a valuable resource for future research in this field.

ACKNOWLEDGMENTS

This research was partially supported by the National Science Foundation under Grant Numbers CNS-1704899, CNS-1815274, CNS-1943396, CNS-1837022, and CMMI-2218809, as well as COGNISENSE, one of seven centers in JUMP 2.0, a Semiconductor Research Corporation (SRC) program sponsored by DARPA. The views and conclusions contained here are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of Columbia University, NSF, SRC, DARPA, or the U.S. Government or any of its agencies.

REFERENCES

[1] 2022. National Household Travel Survey. https://www.nike.com/hr/a/how-long-does-it-take-to-walk-a-mile.
[2] 2023. Con Edison Hosting Capacity Web Application. https://coned.maps.arcgis.com/apps/MapSeries/index.html?appid=edce09020bba4f999c06c462e5458ac7.
[3] 2023. Electric utilities in the U.S. - statistics & facts. https://www.statista.com/topics/2597/electric-utilities/#topicOverview.

[4] 2023. Electric Vehicle Registration Map. https://www.nyserda.ny.gov/All-Programs/chargeny/support-electric/map-of-ev-registrations.

[5] 2023. ENERGY MARKET & OPERATIONAL DATA. https://www.nyiso.com/energy-market-operational-data.

[6] 2023. New York State Energy Research and Development Authority. https://www.nyserda.ny.gov/.

[7] 2023. PlugShare - EV Charging Station Map. https://www.plugshare.com/.

[8] 2023. TLC Trip Record Data. https://www.nyc.gov/site/tlc/about/tlc-trip-record-data.page.

[9] Ali Saadon Al-Ogaili, Tengku Juhana Tengku Hashim, Nur Azzammudin Rahmat, Agileswari K Ramasamy, Marayati Binti Marsadek, Mohammad Faisal, and Mahammad A Hannan. 2019. Review on scheduling, clustering, and forecasting strategies for controlling electric vehicle charging: Challenges and recommendations. *Ieee Access* 7 (2019), 128353–128371.

[10] Muhammad Umair Ali, Amad Zafar, Sarvar Hussain Nengroo, Sadam Hussain, Muhammad Junaid Alvi, and Hee-Je Kim. 2019. Towards a smarter battery management system for electric vehicle applications: A critical review of lithium-ion battery state of charge estimation. *Energies* 12, 3 (2019), 446.

[11] Omar Isaac Asensio, M Cade Lawson, and Camila Z Apablaza. 2021. Electric vehicle charging stations in the workplace with high-resolution data from casual and habitual users. *Scientific Data* 8, 1 (2021), 168.

[12] Yun Bao, Wenbin Dong, and Dian Wang. 2018. Online internal resistance measurement application in lithium ion battery capacity and state of charge estimation. *Energies* 11, 5 (2018), 1073.

[13] Pam Barnhart. 2022. Statistical Data and summaries. https://dmv.ny.gov/about-dmv/statistical-summaries

[14] TomTom International BV. 2023. TomTom. https://www.tomtom.com/. Accessed: 2023-02-08.

[15] Inc ChargePoint. 2023. ChargePoint. https://www.chargepoint.com/. Accessed: 2023-02-08.

[16] Xiaolei Di, Yu Xiao, Chao Zhu, Yang Deng, Qinpei Zhao, and Weixiong Rao. 2019. Traffic congestion prediction by spatiotemporal propagation patterns. In *2019 20th IEEE International Conference on Mobile Data Management (MDM)*. IEEE, 298–303.

[17] Fei Feng, Rengui Lu, and Chunbo Zhu. 2014. A combined state of charge estimation method for lithium-ion batteries used in a wide ambient temperature range. *Energies* 7, 5 (2014), 3004–3032.

[18] Jyotirmoy Gope and Sanjay Kumar Jain. 2017. A survey on solving cold start problem in recommender systems. In *2017 International Conference on Computing, Communication and Automation (ICCCA)*. IEEE, 133–138.

[19] David Hale, Georgios Chrysikopoulos, Alexandra Kondyli, and Amir Ghiasi. 2021. Evaluation of data-driven performance measures for comparing and ranking traffic bottlenecks. *IET Intelligent Transport Systems* 15, 4 (2021), 504–513.

[20] Dickson NT How, MA Hannan, MS Hossain Lipu, and Pin Jern Ker. 2019. State of charge estimation for lithium-ion batteries using model-based and data-driven methods: A review. *Ieee Access* 7 (2019), 136116–136136.

[21] Feliz Kristianto Karnadi, Zhi Hai Mo, and Kun-chan Lan. 2007. Rapid generation of realistic mobility models for VANET. In *2007 IEEE wireless communications and networking conference*. IEEE, 2506–2511.

[22] Seoungbum Kim and Benjamin Coifman. 2014. Comparing INRIX speed data against concurrent loop detector stations over several months. *Transportation Research Part C: Emerging Technologies* 49 (2014), 59–72. https://doi.org/10.1016/j.trc.2014.10.002

[23] Xuan Nhat Lam, Thuc Vu, Trong Duc Le, and Anh Duc Duong. 2008. Addressing cold-start problem in recommendation systems. In *Proceedings of the 2nd international conference on Ubiquitous information management and communication*. 208–211.

[24] Regina Lamedica, Alberto Geri, Fabio Massimo Gatta, Silvia Sangiovanni, Marco Maccioni, and Alessandro Ruvio. 2019. Integrating electric vehicles in microgrids: Overview on hosting capacity and new controls. *IEEE Transactions on Industry Applications* 55, 6 (2019), 7338–7346.

[25] James Larminie and John Lowry. 2012. *Electric vehicle technology explained*. John Wiley & Sons.

[26] Zachary J Lee, Tongxin Li, and Steven H Low. 2019. ACN-data: Analysis and applications of an open EV charging dataset. In *Proceedings of the Tenth ACM International Conference on Future Energy Systems*. 139–149.

[27] Zachary J. Lee, Sunash Sharma, Daniel Johansson, and Steven H. Low. 2021. ACN-Sim: An Open-Source Simulator for Data-Driven Electric Vehicle Charging Research. *IEEE Transactions on Smart Grid* 12, 6 (2021), 5113–5123. https://doi.org/10.1109/TSG.2021.3103156

[28] Pablo Alvarez Lopez, Michael Behrisch, Laura Bieker-Walz, Jakob Erdmann, Yun-Pang Flötteröd, Robert Hilbrich, Leonhard Lücken, Johannes Rummel, Peter Wagner, and Evamarie Wießner. 2018. Microscopic Traffic Simulation using SUMO, In The 21st IEEE International Conference on Intelligent Transportation Systems. *IEEE Intelligent Transportation Systems Conference (ITSC)*. https://elib.dlr.de/124092/

[29] Yunfei Mu, Jianzhong Wu, Nick Jenkins, Hongjie Jia, and Chengshan Wang. 2014. A spatial–temporal model for grid impact analysis of plug-in electric vehicles. *Applied Energy* 114 (2014), 456–465.

[30] Jingping Nie, Lanxiang Hu, Yian Liu, Yuang Fan, Matthias Preindl, and Xiaofan Jiang. 2022. Human-centric data-driven optimization and recommendation in EV-interfaced grid at city scale. In *Proceedings of the 9th ACM International Conference on Systems for Energy-Efficient Buildings, Cities, and Transportation*. 295–296.

[31] Jingping Nie, Yanchen Liu, Liwei Zhou, Xiaofan Jiang, and Matthias Preindl. 2022. Deep Reinforcement Learning Based Approach for Optimal Power Flow of Microgrid with Grid Services Implementation. In *2022 IEEE Transportation Electrification Conference & Expo (ITEC)*. IEEE, 1148–1153.

[32] Jingping Nie, Liwei Zhou, Margaret Frances Kaye, Christine Cecilia Silveira, Afam Nwokolo, Xiaofan Jiang, and Matthias Preindl. 2022. High-Performance Optimal Power Flow Estimation for EV-Interfaced Microgrids With Standardized Grid Services. *IEEE Transactions on Industry Applications* 59, 1 (2022), 1199–1211.

[33] Geunseob Oh, David J Leblanc, and Huei Peng. 2020. Vehicle energy dataset (ved), a large-scale dataset for vehicle energy consumption research. *IEEE Transactions on Intelligent Transportation Systems* 23, 4 (2020), 3302–3312.

[34] Gregory L Plett. 2004. Extended Kalman filtering for battery management systems of LiPB-based HEV battery packs: Part 3. State and parameter estimation. *Journal of Power sources* 134, 2 (2004), 277–292.

[35] Anuj Sharma, Vesal Ahsani, Sandeep Rawat, et al. 2017. Evaluation of opportunities and challenges of using INRIX data for real-time performance monitoring and historical trend assessment. *United States Department of Transportation Report* (2017).

[36] Prashant Shrivastava, Tey Kok Soon, Mohd Yamani Idna Bin Idris, and Saad Mekhilef. 2019. Overview of model-based online state-of-charge estimation using Kalman filter family for lithium-ion batteries. *Renewable and Sustainable Energy Reviews* 113 (2019), 109233.

[37] Åse Lekang Sørensen. 2021. Residential electric vehicle charging datasets from apartment buildings.

[38] Mogile Technologies. 2023. ChargeHub. https://chargehub.com/. Accessed: 2023-02-08.

[39] Jia Yu, Zongsi Zhang, and Mohamed Sarwat. 2019. Spatial Data Management in Apache Spark: The GeoSpark Perspective and Beyond. *Geoinformatica* 23, 1 (jan 2019), 37–78. https://doi.org/10.1007/s10707-018-0330-9

[40] Yanhui Zhang, Wenji Song, Shili Lin, and Ziping Feng. 2014. A novel model of the initial state of charge estimation for LiFePO4 batteries. *Journal of Power Sources* 248 (2014), 1028–1033.

[41] Ou Zheng, Mohamed Abdel-Aty, Lishengsa Yue, Amr Abdelraouf, Zijin Wang, and Nada Mahmoud. 2022. CitySim: A Drone-Based Vehicle Trajectory Dataset for Safety Oriented Research and Digital Twins. *arXiv preprint arXiv:2208.11036* (2022).