

A Holistic Approach for Equity-aware Carbon Reduction of Ridesharing Platforms

Mahsa Sahebdel* Ali Zeynali[†] Noman Bashir[‡] Prashant Shenoy[§]
Mohammad Hajiesmaili[¶]

February 6, 2024

Abstract

In recent years, ridesharing services have revolutionized personal mobility, offering convenient on-demand transportation anytime. While early proponents of ridesharing suggested that these services would reduce the overall carbon emissions of the transportation sector, recent studies reported a type of rebound effect showing substantial carbon emissions of ridesharing platforms, mainly due to their deadhead miles traveled by a ride-share car between two consecutive rides. However, reducing deadhead miles' emissions can incur longer waiting times for riders and starvation of ride assignments for some drivers. Therefore, any efforts towards reducing the carbon emissions from ridesharing platforms must consider the impact on the quality of service, e.g., waiting time, and on the fair and equitable distribution of rides across drivers.

This paper proposes a holistic approach to reduce the carbon emissions of ridesharing platforms while minimizing the degradation in user waiting times and equitable ride assignments across drivers. Towards this end, we decompose the global carbon reduction problem into two related sub-problems: carbon- and equity-aware ride assignment and fuel-efficient routing. For the ride assignment problem, we consider the trade-off between the amount of carbon reduction and the rider's waiting time and propose simple yet efficient algorithms to handle the conflicting trade-offs. For the routing problem, we analyze the impact of fuel-efficient routing in reducing the carbon footprint, trip duration, and driver efficiency of ridesharing platforms using route data from Google Maps. Our comprehensive trace-driven experimental results show substantial emissions reduction of our proposed algorithms with only a graceful increase in riders' waiting times. Finally, we release "E²-RideKit", a toolkit that allows researchers to augment ridesharing datasets with emissions and equity information, enabling further research on emissions analysis and improvement of ridesharing platforms.

*University of Massachusetts Amherst. Email: msahebdelala@umass.edu.

[†]University of Massachusetts Amherst. Email: azeynali@umass.edu.

[‡]University of Massachusetts Amherst. Email: nbashir@umass.edu.

[§]University of Massachusetts Amherst. Email: shenoy@cs.umass.edu.

[¶]University of Massachusetts Amherst. Email: hajiesmaili@cs.umass.edu.

1 Introduction

In 2021, the United States accounted for about 17% of global greenhouse gas (GHG) emissions. Of the 4.6 Billion Metric Tons (BMT) energy-related emissions produced that year, 38% was from the transportation sector—the highest share among six sectors [1]. Globally, transportation has a similar impact, contributing 37% of CO₂ emissions in 2021 [2]. In particular, urban mobility is estimated to contribute 40% of CO₂ emissions from road transport [3], and given urbanization trends, demand is projected to more than double by 2050 [4].

The impact of ridesharing services on the transportation market has grown quickly during the past few years as a sustainable alternative to individual vehicle ownership. As reported by Bloomberg, the size of ridesharing services was valued at \$69.3B in 2022 and is expected to hit \$205.8B at the end of 2030 [5]. These services have revolutionized people’s travel by providing convenient access to individual or shared vehicles based on their requested pick-up and drop-off locations. Consequently, ridesharing services, such as Uber, Lyft, Grab, and Didi, have become immensely successful due to their promise of personal on-demand mobility at any time [6].

Early proponents of ridesharing suggested that these services would reduce reliance on privately-owned cars, reduce traffic congestion, and reduce carbon emissions, with early studies estimating that at least five private vehicles would be replaced for each shared car and there would likely be carbon emission reductions if shared cars were newer vehicles [7]. However, the success of these services has resulted in an increase in traffic and more congestion on roads—a rebound effect [8]. For example, in New York City, ridesharing has been shown to constitute more than 50% of road traffic [9, 10]. Another study has estimated that a typical ridesharing trip is less efficient than a personal car trip, mainly due to “deadhead” miles traveled by a ride-share vehicle between consecutive hired rides, and that this generates 47% more CO₂ emissions than an equivalent private car ride [11, 12]. The study also showed that the greater convenience of ridesharing has steered passengers away from public transit options. Others have reported that the deadhead miles of ridesharing services account for an estimated 36-45% of their overall distance traveled [13, 14, 15].

The prevalence of deadhead miles and resulting carbon emissions presents new opportunities to develop emission-aware ride assignment algorithms. However, any carbon-aware optimization of ridesharing systems must consider its impact on riders and drivers. Prior work shows that existing ride assignment algorithms may be unfair to drivers from certain demographics based on their gender, age, and race [16, 17, 18]. An emission-aware ride assignment can perpetuate such disparities by prioritizing drivers with newer, more fuel-efficient vehicles. At the same time, new emission-aware ride assignment algorithms should not increase rider’s waiting times, which can negatively impact user satisfaction. Therefore, any efforts towards reducing the carbon emissions from ridesharing platforms should carefully consider its impact on the quality of service (QoS), e.g., waiting time, and equity, e.g., distribution of rides across drivers. While there has been significant work on optimizing various facets of ridesharing systems [19, 20, 21], the problem of emission-aware ride assignment to green ridesharing systems has not seen much research attention.

The design space for emission-aware ridesharing optimization is two-fold: (1) ride assignment to reduce the emissions from the deadhead miles between two trips; (2) the routing strategy of a trip to optimize emissions instead of other factors such as trip duration or distance. In both problems, there is a trade-off between emissions, QoS, and equity implications. A ride assignment based on a single objective of emissions reduction may degrade both the *rider* and *driver*’s QoS. For example, emission-aware ride assignments may increase the rider’s waiting time by picking a vehicle from a soon-to-be-completed trip to minimize the deadhead miles. Similarly, during trip routing, the extended time in picking the most fuel-efficient route instead of the fastest route may degrade user waiting times. The equity impacts of emission-aware ride assignments are more subtle. The emission-aware ride

assignment may prefer electric or low-emissions vehicles, which is unfair to drivers with high-emissions vehicles that tend to belong to low-income communities. Alternatively, the emissions-aware ride assignment policy may assign trips with inevitably higher deadhead miles to low-emissions vehicles, which then decreases the efficiency of their service due to a higher deadhead-to-trip ratio.

In this paper, we take a holistic approach to emission-aware ride assignment and trip routing in ridesharing platforms that considers both the QoS and equity implications. We analyze the performance and the trade-offs between their emissions reduction and QoS of riders and drivers using a real-world ridesharing dataset. In doing so, our paper makes the following contributions:

1. **Emission-aware ride assignment:** We develop emissions-aware ride assignment algorithms aimed at minimizing emissions from deadhead miles while maintaining the quality of service (QoS) for both riders and drivers. We first introduce an offline solution that showcases the substantial potential for emission reduction through an emission-aware ride assignment policy. Following that, we introduce an online strategy that allows for a configurable trade-off between emission reductions and QoS, as measured by rider waiting times.
2. **Trip routing:** In addition to the ride assignment strategy, the routing policy also plays a significant role in determining overall carbon emissions. We utilize route data provided by popular navigation apps such as Google Maps to evaluate the trade-offs among time (the fastest route), distance (the shortest route), and fuel consumption (the greenest route) in routing following emission-aware ride assignments.
3. **E²-RideKit design:** As an independent contribution, we present E²-RideKit, which augments an arbitrary ridesharing dataset with per-ride carbon emission and per-driver/passenger socioeconomic demographic information. It allows users to configure EV penetration in the dataset, plugin complementary models, and extend the toolkit to add additional information. As an example use case, we augment the widely-used RideAustin dataset [22] with carbon emission and equity information. In doing so, we outline the challenges we faced in augmenting the dataset and how we solved them. We hypothesize that E²-RideKit enables researchers to augment ridesharing datasets with emissions and equity information, enabling new ridesharing analytics and optimizations.
4. **Experimental results:** Finally, leveraging our toolkit, we conduct a comprehensive evaluation of our emission-aware ride assignment algorithm and trip routing policies using the RideAustin dataset to produce multiple key findings. **First**, we demonstrate that our emission-aware ride assignment algorithms can decrease emissions from deadhead miles by up to 60% with a mere 4% increase in waiting time for the riders. **Second**, replacing just 5% of the current fleet with EVs can improve deadhead miles emission reduction to 67%. **Third**, optimizing trip routing for emissions does not result in favorable tradeoffs; emissions can only be reduced by 4.2% w.r.t. the fastest route while increasing the trip duration by 3%. **Finally**, combining emission-aware ride assignment and trip routing can reduce the fleet-level emissions by 26%, at the expense of a 6.2% increase in trip duration. The emission reduction can be improved to 29% by replacing only 5% of the current fleet with EVs.

2 Problem Statement and Challenges

Addressing the general problem of emissions reduction in ridesharing¹platforms involves two main challenges. First, the routing of trips within the platform is important. Routing algorithms can be tailored not only for the shortest distance or fastest route but also to minimize fuel consumption and

¹Our work does not address carpooling or ridesharing of passengers.

emissions. Second, optimizing ride assignments to reduce emissions is also crucial. This involves efficiently allocating vehicle rides to minimize deadhead mile between the current trip’s drop-off location and the next pick-up location. Achieving this goal is complex because it requires balancing reducing emissions and maintaining a high-quality experience for drivers and riders. These two interconnected components of the problem require careful consideration and trade-offs to ensure a sustainable and equitable ridesharing ecosystem that simultaneously addresses environmental and service quality concerns.

The fuel-efficient routing problem. Routing algorithms are pivotal in determining the optimal route during a trip. These algorithms consider various factors, including travel time, distance, and emissions produced. Variables such as route distance, traffic conditions, and road quality influence travel time and emissions associated with a particular route. It is important to note that the route with the lowest emissions may not always align with the shortest distance or the fastest travel time. As a result, different routing algorithms may yield varying emissions and waiting times results.

The emission-aware ride assignment. The ride assignment strategy also impacts the emissions due to the deadhead mile of a trip. Since drivers are located at varying distances from passengers, and their vehicles have distinct unit-distance emissions (emissions produced per unit of distance traveled), the passenger-to-driver assignment algorithm significantly impacts the overall emissions of the ridesharing system and the average waiting time for passengers. While some simplistic assignment algorithms may reduce emissions for individual trips by assigning passengers to drivers with the lowest emissions, this approach can lead to substantially extended waiting times if the low-emission driver is situated far from the passenger, potentially degrading the passenger’s satisfaction. Furthermore, assigning passengers to drivers with the lowest emissions does not inherently result in long-term optimal ride assignments for the entire ridesharing service. On the other hand, a strategy that prioritizes assigning passengers to the nearest available drivers aims to minimize waiting times but may inadvertently result in higher carbon emissions, as it overlooks variations in unit-distance emissions produced by different vehicles in the fleet.

An illustrative example. In Figure 1, we illustrate the benefits of our system-level emissions-aware ride assignment over per-ride optimization strategies [23]. Let’s consider a scenario where a driver M_1 carries a passenger at time t_0 , and the passenger will be dropped off at t_5 . At t_3 , two passengers request rides. In a per-ride optimization, the driver M_1 is assigned to passenger N_2 due to its proximity, and the passenger N_3 is assigned to the driver M_2 , yielding a total of 61 deadhead miles. However, if we adopt a system-level view of the ride assignment, the driver M_1 would be assigned to the passenger N_3 , and driver M_2 to passenger N_2 , even though this is not the best assignment for driver M_1 . This results in 40 deadhead miles, significantly reducing the deadhead mile emissions.

Problem formulation. At a high level, the ride-assignment problem is an extended version of the classic online bipartite matching problem where two sets of disjoint nodes represent riders and drivers. However, modeling the topology constraints, i.e., determining the set of available drivers and their corresponding assignment cost for a new ride request, is highly challenging since other rides might already occupy some drivers, where some of those drivers might be available soon since they are close to finishing their current ride. A rigorous formulation must carefully capture multiple complex topology-constrained and time-coupled relationships between riders and drivers. We formulate the emission-aware one-on-one ride assignments within a ridesharing platform by considering a high-level abstraction of the topology constraint of the problem in an online setting. We consider a ridesharing platform comprising M drivers. Over a certain time horizon, N ride requests are generated. The objective of the ride assignment algorithm is to assign an available empty car to each request such that it minimizes the long-term emissions from the total distance traveled. This encompasses rides with passengers and the deadhead mile between dropping off one passenger and picking up the next.

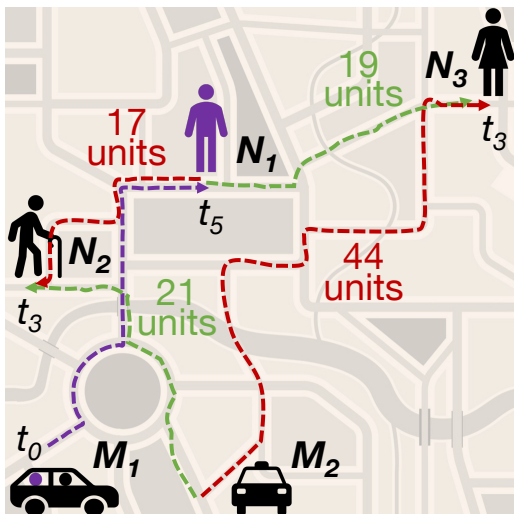


Figure 1: *Ride assignment with the goal of minimizing deadhead miles yields lower such miles than per-ride optimization (40 vs. 61 units).*

We describe a simplified version of the Emission-aware Ride Assignment Problem (ERAP) as follows.

$$\begin{aligned}
 \text{[ERAP]} \quad & \min \sum_{n=1}^N \sum_{m=1}^{M_n} (e_t(n, m) + e_d(n, m)) x_{n,m} \\
 & \text{s.t., } \sum_{m=1}^{M_n} x_{n,m} = 1, \quad n \in [N], \\
 & \text{vars., } x_{n,m} = \{0, 1\}, \quad n \in [N], m \in [M_n],
 \end{aligned}$$

where $e_t(n, m)$ are the trip emissions due to the assignment of passenger n to driver m and $e_d(n, m)$ represents the emissions due to the deadhead mile for driver m to pickup passenger n . The routing algorithms significantly impact these two terms. Also, $x_{n,m}$ is derived from a ride assignment algorithm and is a binary optimization variable where $x_{n,m} = 1$ if m is assigned to n ; 0, otherwise.

Balancing the maximum potential of emissions reduction and riders' QoS. In ERAP, parameter M_n denotes the set of available cars for passenger n that could include both currently available and soon-to-be-completed cars for ride n . Hence, M_n is a crucial parameter that balances the maximum potential of emissions reduction and riders' waiting time (QoS), i.e., the more willingness to increase the waiting time of ride n , the more cars will be eligible to be included in M_n . In our algorithm in Section 3.2, we define a threshold parameter ϕ that determines the set of available cars for a new ride, and by adjusting parameter ϕ , the set of available cars M_n changes; hence, the algorithm could be tuned to achieve a desired trade-offs between emission reduction and user waiting time.

Lastly, we note that one could implicitly leverage parameter M_n to impose driver's equity constraints. That is, there could be another separate module that records the equity-related metrics for drivers, e.g., deadhead-to-trip ratio and the number of assigned trips, and then based on the equity status of each driver, the available cars for the new ride might be determined, e.g., the driver equity module may exclude the drivers with a large number of previously assigned rides to make the ride assignment equitable for other available drivers. That said, we emphasize that our current modeling does not *explicitly* model the optimization of equity from the drivers' perspective. Designing equitable ride assignments for drivers is a significant further direction of our work.

Algorithm 1: ERA(N, M)

```
1 Initialize  $Q_i = \{\}$ , for  $i \in [1, 2, \dots, M]$ ;  
2 initial_assignment =  $\{Q_i \mid \forall i \in [1, 2, \dots, M]\}$ ;  
3  $\mathcal{A} = \{\text{initial\_assignment}\}$ ;  
4 for  $n \in [1, 2, \dots, N]$  do  
5    $\mathcal{A}' = \{\}$ ;  
6   for  $a \in \mathcal{A}$  do  
7     children = Children( $a, n$ );  
8      $\mathcal{A}' \leftarrow \mathcal{A}' \cup \text{children}$ ;  
9   emissionmin = Minimum emission of assignments in  $\mathcal{A}'$ ;  
10   $\mathcal{A} \leftarrow \{a \mid a \in \mathcal{A}', \text{emission}_h(a) \leq \text{emission}_{\text{min}}\}$ 
```

3 Ride Assignment Algorithms

In Section 3.1, we first present an offline ride assignment algorithm, **ERA**, designed to assign rides to minimize total emissions near optimally. While **ERA** is not a practical algorithm, it can show the maximum potential of emissions reduction with the sole objective of reducing the emissions. In Section 3.2, we then present **TORA**, an online emission-aware ride assignment algorithm that balances total emissions and the average passenger waiting time.

3.1 ERA: An Offline Emission-aware Ride Assignment Algorithm

ERA is an offline algorithm (Algorithm 1) designed to assign a set of N riders to a fleet of M drivers with the primary objective of minimizing emissions. The algorithm initialization phase (Lines 1-3) starts with an empty assignment for all M drivers. Subsequently, the algorithm sequentially considers each passenger n and generates all possible assignments of that passenger to the M available drivers (Line 7). Importantly, the new assignments do not modify assignments to previous passengers.

The algorithm relies on a heuristic function denoted as $\text{emission}_h(a)$ to estimate the emission associated with a specific assignment and calculates the minimum emissions in assigning all passengers to the drivers. Specifically, for any assignment $n \rightarrow m$ in the partial assignment a , the passenger n is obligatorily assigned to driver m . We conduct an extensive numerical analysis **ERA** to define a heuristic function, $\text{emission}_h(a)$. It *estimates* the emissions generated by the system based on the assignments made in a partial assignment a . The function estimates the emissions by calculating the deadhead distance for each driver. The deadhead distance is the distance between the drop-off location of a passenger and the pick-up location of the next passenger assigned to the same driver. In addition, the algorithm takes into account the lower bound for the deadhead distance of any unassigned passenger n . This lower bound is calculated by finding the minimum distance required to drive to the pick-up location of passenger n from the drop-off location of any passenger that arrived before n in the assignment sequence. By deriving the lower bound on the minimum required deadhead distance for each unassigned passenger and leveraging the known emission rate per mile for each driver, the algorithm can effectively estimate the minimum emission for each ride. The total estimated emission is calculated as the sum of the emissions associated with the assigned passengers and the minimum emissions of the unassigned passengers, all for the partial assignment a .

Finally, the algorithm keeps the assignments with the lowest estimated emission and starts branching from those (Lines 10-11). In this algorithm, Children(assignment _{$n-1$} , n) returns a set of possible assignments of passenger index n to drivers without changing any previously assigned passengers, assignment _{$n-1$} , (Line 7). In Figure 2, we present emissions from the deadhead miles and passenger waiting times for the current assignments in the dataset and the new assignments based on our emission-aware offline algorithm.

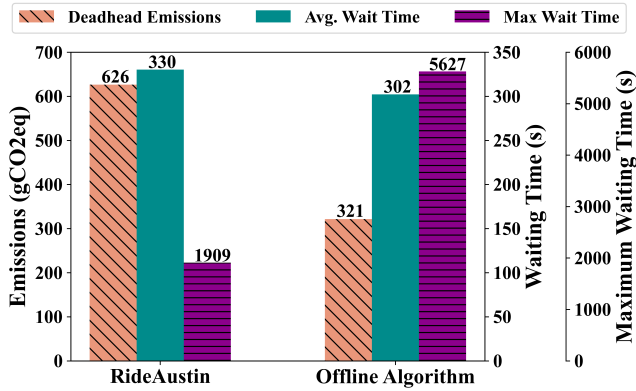


Figure 2: *Opportunity analysis: Comparison of emissions from deadhead miles and waiting time for the default ride assignment and our offline emission-aware ride assignment.*

We note that the computational complexity of the ERA algorithm is substantial, rendering it impractical to seek near-optimal solutions for datasets with hundred thousands of trips. Consequently, for this preliminary feasibility study, we have limited our experimentation to a curated dataset comprising 142 samples. We randomly sampled 142 trips, completed by 14 unique drivers, from the RideAustin dataset [22] (detailed on Section 6.1) on December 2, 2016. The left y -axis presents the total carbon emissions, measured in grams of carbon dioxide equivalent (g.CO2eq), for the deadhead miles of all the trips². The y -axes on the right show the waiting time (seconds), measured as the time period between a passenger posting the request and the driver picking up at different scales. Our results demonstrate that our proposed ERA algorithm, albeit offline with complete knowledge of future rides, can reduce the deadhead miles emissions by 48.7% (from 626gCO2eq to 321gCO2eq). Importantly, the average waiting across all trips also decreased by 8.5% (from 330s to 302s) as ERA reduced the deadhead miles. Still, the longest waiting time increased by $2.94\times$ to almost 94 minutes from the 32 minutes observed for the default ride assignment.

The above initial results show the potential of carbon emissions reduction by changing the assignment objective. However, ERA is an offline algorithm; hence, it is not practically a relevant choice. Hence, in what follows, we present TORA as an online and practical algorithm for emission-aware ride assignment.

3.2 TORA: An Online Threshold-based Ride Assignment Algorithm

TORA is an online ride-assignment algorithm that controls the trade-off between waiting time and vehicle carbon emissions. To minimize waiting time for the passenger n , TORA first finds the closest available driver to the passenger and then compares the distance and deadhead emission produced by other available drivers with the closest driver. TORA calculates the *Emission-to-Distance* (E2D) ratio for every driver and selects the appropriate driver for the passenger based on E2D values. More specifically, E2D is defined as the ratio between the difference of the deadhead emission of a driver and the deadhead emission of the closest driver over the difference of the distance of the passenger to the two drivers, i.e.,

$$\text{E2D}(m, c) := \frac{e_d(n, c) - e_d(n, m)}{\text{distance}(m) - \text{distance}(c)}, \quad (1)$$

where $\text{distance}(c)$, and $\text{distance}(m)$ show the distance of passenger n to the closest driver c and driver m .

²This measurement is derived from multiplying the miles traveled by the fuel efficiency of the vehicle, denoted in grams of CO2 emitted per kilometer (gCO2/km). See the implementation section for additional details.

Algorithm 2: TORA(n, ϕ)

```
1  $c \leftarrow$  closest available driver to  $n$ ;  
2  $m \leftarrow \arg \max_{m'} \text{E2D}(m', c)$ ;  
3 if  $\phi E_0 < \text{E2D}(m, c)$  then  
4    $\lfloor$  return  $m$ ;  
5 else  
6    $\lfloor$  return  $c$ ;
```

The pseudocode for TORA is outlined in Algorithm 2. It leverages a parameter, ϕ , to balance the trade-off between reducing passenger waiting time and minimizing vehicle deadhead emissions. Initially, TORA identifies the closest available driver and computes E2D value for other available drivers (Lines 1-2). Next, it picks a driver with the highest E2D ratio such that $\text{E2D} > \phi \times E_0$, where E_0 are the emissions for a baseline vehicle over a unit distance trip. If such a driver exists, TORA selects and returns it (Lines 6-7). Otherwise, it defaults to assigning the closest available driver.

The underlying concept of this algorithm is straightforward: it replaces the nearest driver with a more distant one if the latter emits ϕE_0 fewer emissions per additional unit distance traveled. The lower values of ϕ prioritize drivers with lower emissions, whereas higher values prioritize selecting the closest driver, potentially increasing emissions. To better understand how different values of ϕ impact the passenger’s waiting time, we can calculate the upper bound on the distance of driver m to the passenger when TORA is willing to replace the closest driver c by m . From the definition of $\text{E2D}(m, c)$ in Equation 1, TORA replaces closest driver c by driver m if

$$\phi E_0 < \frac{e_c \cdot \text{distance}(c) - e_m \cdot \text{distance}(m)}{\text{distance}(m) - \text{distance}(c)}, \quad (2)$$

where e_c and e_m denote the unit-distance emission for the driver c and m . By rearranging Eq. (2), TORA replace driver c with driver m if

$$\frac{\text{distance}(m)}{\text{distance}(c)} < \frac{e_c/E_0 + \phi}{e_m/E_0 + \phi}, \quad (3)$$

where the ratio of $\text{distance}(m)/\text{distance}(c)$ approximates the ratio of waiting time for the passenger with driver m over driver c . Eq. (3) shows that TORA limits the distance of the lower emission driver based on its and closest driver’s unit-distance emissions.

4 Route Choice Trade-Offs

In this section, we study the impact of different route options on the total distance, trip duration, and emission produced during the trip. We consider the shortest route, fastest route, and most fuel-efficient route as the three possible choices of driver to travel from pick-up to drop-off location of passenger. We collected the data for the three route options — distance, duration, and fuel consumption — for 44,794 trips using the Google Maps API [24]. In our analysis, we categorize trips into short- (< 1 mile), medium- (1–10 miles), and long-distance (> 10 miles) trips (based on the length of shortest route) to emphasize the impact of different routing options on different trips. Our three route options also give us three metrics for our analysis: distance traveled, time taken, and emissions produced. We choose one route option as a baseline (optimize one metric) and evaluate how worse off the other two routes are based on the baseline metric for all trip categories.

In Figure 3(a), we examine the percentage increase in distance for the fastest and fuel-efficient route options when compared to the baseline shortest distance route (y -axis) for the three trip categories (x -axis). The fastest route takes a longer route, probably to avoid traffic congestion, by up to

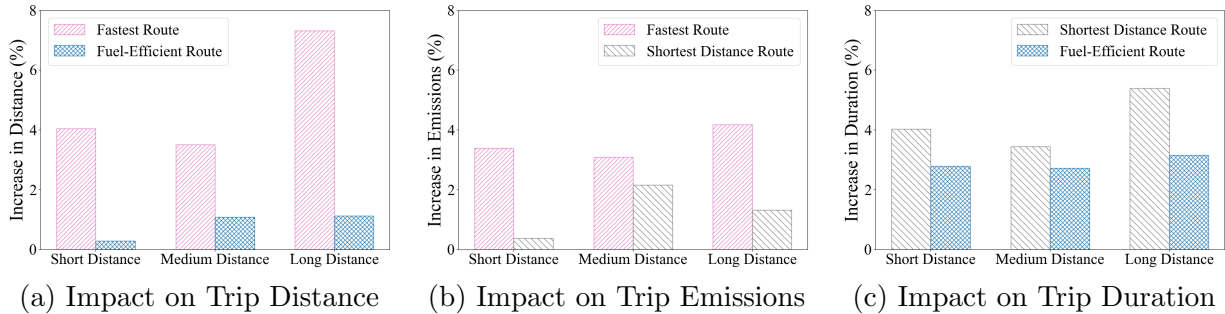


Figure 3: Average increase in distance (a), emission (b), and duration (c) for different routing algorithms to the optimal routes for short, medium, and long-distance trips for three trip categories in our analysis.

7.5%, depending on the length of the trip. The fuel-efficient route needs to balance the savings from traveling the shortest distance and the cost of idling in a traffic jam; it drives longer highly judiciously. Figure 3(b) shows the increase in trip emissions for the shortest distance and fastest routes. The increase in emissions is smaller than the increase in distance we observed in Fig. 3(b). Also, the actual magnitude of increase even for the fastest trip is around 4%. The differences may further diminish in practice, especially in cities where traffic patterns change quickly and unexpectedly. Finally, in Figure 3(c) we evaluate the time cost of the shortest distance and fuel-efficient route options. As we observed before, shortest distance routes are most likely congested, and choosing them would cost up to 6% more time than the shortest duration route. Since the fuel-efficient route option travels longer distances but avoids congestion, its increase in time is smaller.

Our analysis showed that the choice of route had varying impacts on trip duration, distance, and emissions, with the shortest, fastest, and fuel-efficient routes exhibiting different trade-offs.

Key takeaways. *Our analysis yields a few interesting and surprising insights. It shows that choosing the fuel-efficient route may not always yield huge fuel savings. However, since the distance and time increase for the fuel-efficient trip are minute, <1% and <2.5%, respectively, the small savings do not come at a huge cost. Also, the results are subject to local road networks, traffic conditions, and driver habits.*

5 E²-RideKit Toolkit

Our work on developing emission-aware ride assignment and routing algorithms requires ridesharing datasets to include carbon emission information for the deadhead miles and individual trips, and equity information on the drivers and riders. However, no existing ridesharing datasets contain such information. Augmenting the existing datasets with emission and equity information requires solving multiple challenges. In this section, we present the design and implementation E²-RideKit³ for augmenting ridesharing datasets as an independent contribution. To make the toolkit complete and independently useful, we implement several additional functionalities not used directly for our algorithmic work in this paper. However, they are of independent interest for emissions research in the ridesharing ecosystems.

³https://github.com/Mahsahebdel/e2_ridekit

5.1 E²-RideKit Design

The design of E²-RideKit needs to solve three key challenges: (1) building vehicle- and congestion-aware emission models for internal combustion engine (ICE) powered vehicles, (2) modeling congestion-aware discharging and location-aware charging models for Electric Vehicles (EVs), and (3) defining, quantifying, and adding missing information on equity-related metrics. We solve these challenges as a part of our design, which we describe next.

1. ARCHITECTURE. The high-level architecture of E²-RideKit consists of three major modules, as shown in Figure 4. The three major modules are (i) **Ride Enhancer** module, (ii) **Emission Estimator** module, and (iii) **Equity Embedder** module.

Ride Enhancer. The goal of this module is to augment the per-ride/trip information from an existing dataset with missing information, if any. For example, a dataset may only provide the pickup and drop-off coordinates without sharing information on the actual route taken by the driver. Since we need to estimate per-ride emissions, which depend on the route the driver takes, **Ride Enhancer** module can leverage a navigation API to generate trajectories. Similarly, the congestion model inside the **Ride Enhancer** module is responsible for estimating congestion during the trip since it is one of the key factors that affect the fuel efficiency of vehicles, which impacts the carbon emissions for a given trip. Note that **Ride Enhancer** module operates on all rides and does not differentiate between rides from an ICE-powered vehicle and an EV.

Emission Estimator. This module is the heart of the E²-RideKit as it estimates the carbon and other GHG emissions for various aspects of each ride. For example, it will estimate the emissions during the actual trip with the passenger and emissions from deadhead miles that the driver had to undertake to pick up a new customer. Since the car make/models, fuel types, and efficiencies of ICE-powered vehicles and EVs are significantly different, we use separate modules for estimating their emissions. Based on the vehicle type, EV or ICE-powered, we use data from either carbon information services such as electricityMaps [25] and Watttime [26] or existing vehicle emissions datasets [27].

Equity Embedder. An equally important and the most challenging module is the **Equity Embedder** module that embeds socioeconomic and demographic information for the drivers and the riders. Due to privacy concerns, dataset collectors and maintainers do not release information that can reveal the identity of the riders or passengers. As a result, this information needs to be estimated at a coarse granularity using publicly available information using resources such as US Census Bureau [28]. Specifically, we use the Geocode API [29] to collect demographic data such as race, population, and median income. The data is provided by the U.S. Census Bureau and is made available by census block, i.e., the smallest geographical unit for which the U.S. Census Bureau collects and provides statistical data. In some cases, the data may not be for the same year as the ride-sharing dataset. For example, for many locations in the US, 2020 is the most recent year with available census data [30]. However, this should not be a significant issue as the demographics of cities, or even blocks, only change over the course of many years. Using per-ride information, our module optionally estimates coarse information on which segment of the city the driver or rider lives. This information can be gleaned from past datasets for the regular drivers and riders using methods such as K-means clustering [31]. We release this module with a disclaimer about the potential inaccuracies in the embedded information.

2. WORKFLOW. We next describe the workflow of E²-RideKit when using it to augment an existing ridesharing dataset. We first split the input dataset into two segments that contain information on individual rides and the other segment containing information on all the drivers and riders if available.

If the dataset does not provide trajectory information, we leverage navigation APIs with generative modeling to estimate trajectories for each ride, i.e., the route between pickup location and destination.

This trajectory information is needed to estimate the trip emissions based on the vehicle used, route taken, and distance traveled. We next adjust the trip emissions based on the estimated travel time index (TTI), which compares the travel time of a given trip with the typical time for the same trip and represents a good proxy for traffic congestion. By leveraging the appropriate emission module, this step handles both ICE-powered vehicles and EVs.

To solve the second challenge, when considering discharging, we will leverage the same approach as the first challenge. To model the charging patterns and constraints, we will leverage publicly-available information on the charging station locations [32], the ratings of the available chargers [32], and the carbon intensity of electricity supply using services such as electricityMaps [25] and Watttime [26]. To solve the final challenge, we will create a meta-dataset that provides the distributions of equity-related metrics augmented with additional information on socio-economic factors from U.S. Census Bureau [28]. Table 1 lists the additional per-trip and meta fields that E²-RideKit add to our example dataset.

5.2 E²-RideKit Implementation

In this section, we describe the implementation of E²-RideKit and the models and datasets we use for our case study, as explained in the following. We combine multiple additional datasets as a part of our implementation and describe them next.

1. *Car emissions dataset.* To model emissions from the ICE-powered vehicles in the dataset, we use the Canada Car Emissions dataset [27]. This extensive and reliable resource provides a comprehensive overview of CO₂ emissions data for various vehicles. This dataset covers a period of 7 years and is sourced from the official open data website of the Canadian government. It consists of detailed information on different aspects and characteristics of 26,075 unique vehicles. The dataset primarily focuses on essential factors such as vehicle’s make, model, year, and carbon emissions, which are fundamental in understanding CO₂ emissions. However, it goes beyond these key attributes and includes valuable information such as fuel type, engine size, transmission type, and other relevant features. Easy access to these additional details through the toolkit will enable researchers and analysts to gain deeper insights into the relationship between car attributes and carbon emissions, facilitating a more comprehensive dataset analysis.

2. *Carbon emissions dataset.* ElectricityMaps [25] and Watttime [26] are valuable services that provide essential information regarding the carbon intensity of electricity supply across different locations. This information is presented as average and marginal hourly values, which can vary depending on the specific location and time. The availability of this data is particularly beneficial for modeling the carbon intensity of electricity supply during the charging of EVs. To accurately model the charging patterns and constraints for EVs, publicly available data on charging station locations and charger ratings are utilized. This data aids in estimating the charging patterns of EVs and determining the time required for a vehicle to be fully charged. However, it is not always feasible to determine the exact time at which an EV is charged. In such cases, a proxy approach is employed, where the daily average carbon intensity value of the electricity supply at the location where the vehicle is charged is utilized to estimate the charging time. By adopting this approach, an estimation of the carbon intensity of the electricity supply at the time of charging is derived and incorporated into the analysis. By leveraging the information from electricityMaps, WattTime, and the data on charging station locations and charger ratings, researchers and analysts can enhance their analysis of the environmental impact of EVs and their charging patterns. This integration allows for a holistic understanding of the relationship between EV charging, the carbon intensity of electricity, and its implications for sustainability and environmental considerations.

3. *Traffic congestion model.* The traffic congestion model is a demand modification and traffic simu-

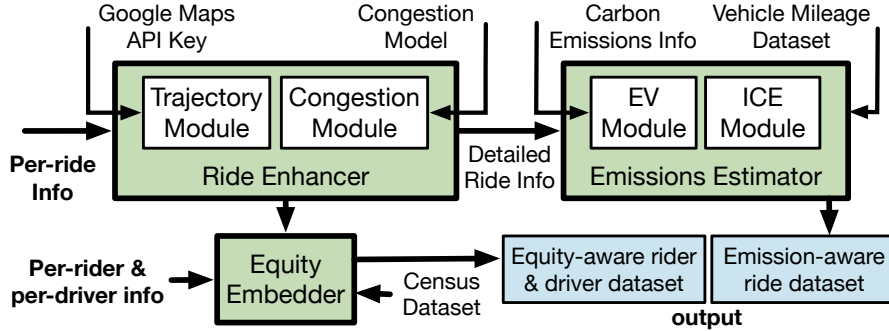


Figure 4: *The high-level architecture, various components, and workflow of E^2 -RideKit.*

Table 1: Additional Fields Added to the RideAustin Dataset.

Trip related fields	fuel efficiency	emissions (deadhead & trip)	congestion index
	EV (yes/no)	demographics (rider & driver)	
Meta fields	EV trajectories	electricity carbon intensity	charger location
	charger type	charging/discharging model	

lation process. The model aims to estimate changes in travel demand and traffic patterns. We next detail the demand modification and traffic simulation processes.

In the demand modification process, we adjust the Origin-Destination (O-D) matrix to reflect changes in travel demand resulting from various factors. This requires data-derived inputs, such as trip proportions associated with specific Traffic Analysis Zones (TAZ) and the proportion of dedicated tours for specific purposes. Users can also set additional simulation parameters, like the penetration and substitution rates. To incorporate changes in travel patterns, we use a Monte Carlo simulation, which accounts for uncertainties in determining trip categories and subsequent destinations by generating multiple potential scenarios. After identifying the affected trips, the OD matrix is modified to reflect the changes introduced in the simulation.

The traffic simulation process involves using the Vehicle Hours Traveled (VHT) metric to understand and quantify the impact of transportation activities on congestion. VHT measures the total travel time of vehicles on the road network during a specific period, considering both the number of vehicles and the duration of their trips. To calculate VHT, we use Static Traffic Assignment, which assigns trips from an Origin-Destination (OD) matrix to the transportation network, considering travel demand, road capacities, and traffic conditions. The goal is to distribute the trips across the network in a way that reflects realistic travel patterns and estimates travel times for the assigned trips. The resulting network data provides valuable information on updated link speeds, travel times, and vehicle volumes, enabling the calculation of VHT by summing up the travel times on each link. By comparing VHT values and analyzing different scenarios, researchers can assess the impact of transportation activities on congestion, aiding in the understanding of traffic dynamics, evaluation of interventions, and development of strategies to mitigate congestion-related issues.

6 Experimental Evaluation

This section provides a comprehensive experimental study to evaluate the deadhead, total emission, and waiting time of a ridesharing platform under ride assignment of TORA with different ϕ values and different routing algorithms explained in Section 4. The key questions for the evaluations and our findings are outlined below.

- Q1 *How does TORA impact the deadhead emissions and waiting times when the routing algorithm remains unchanged?* We find that TORA can improve the deadhead emission significantly with a cost of slightly increasing the waiting time (Key takeaway 1).
- Q2 *How does TORA affect the equity among drivers with varying vehicle emissions?* Our analysis reveals that reduction in emissions requires sacrificing equity by assigning more rides to low-emission vehicles (Key takeaway 2).
- Q3 *How do the TORA and routing algorithms compare regarding their impact on the emission of ridesharing systems?* We find that the impact of the ride assignment algorithm on the emission and waiting time of the ridesharing system is substantially more than the impact of routing algorithms (Key takeaway 3).

6.1 Experimental Setup

We outline the characteristics of the ridesharing dataset we use, various parameters that we vary in our experiments, and the metrics we use to evaluate the efficacy of our approach in reducing emissions and ensuring equity.

Ridesharing dataset. RideAustin, a non-profit ridesharing company based in Austin, Texas, has released a dataset encompassing around 1.5 million trips spanning a 10-month period across 2016 and 2017 [22]. During this period, no other major ridesharing platforms like Uber or Lyft operated in Austin. The dataset contains comprehensive trip-specific information, including the trip’s start and end times, geographical coordinates for both pick-up and drop-off locations, the vehicle’s make and model, distance covered before, during, and after each trip, and unique identification numbers for the driver and passenger.

In our experiments, we use RideAustin dataset consisting of trips between December 1, 2016, and December 10, 2016. This subset comprises 44,794 ridesharing trips conducted by 1,406 drivers. We augment this dataset with emissions information using E²-RideKit, and the additional dataset explained in Section 5. We categorized the vehicles emitting less than 135 g.CO₂eq/km as low-emission vehicles (LEVs), and created three more variants of the dataset by randomly selecting 5%, 10%, and 20% of the vehicles which are not categorized as a low-emission and converting them into electric vehicles (EVs), yielding three additional datasets with 10%, 15%, and 25% low-emission vehicles. During our experiments, we set the unit distance emission equivalent of EVs to be 63.35 g.CO₂eq/km⁴.

Baseline strategies and parameter ranges. For our ride assignment algorithms, we use two baseline ride assignment strategies: 1) default assignment as presented in the RideAustin dataset and 2) shortest distance assignment that assigns the closest driver to the passenger. The second assignment strategy is emulated as an extreme case of TORA when ϕ approaches infinity. We use three variants of TORA with ϕ values of 0.1 (representing an emissions-prioritized assignment), 1 (roughly representing emissions- and -waiting-time-aware assignment), and 7.5 (representing a waiting-time-prioritized assignment) for our ride assignment strategies. For routing, we used three routing strategies: the shortest, the fastest, and the fuel-efficient, as explained in Section 4.

For our experimental parameters, we vary the values of ϕ between 0.001 and 18. We set the value of E_0 to 63.35 g.CO₂eq/km, which is equal to the unit-distance carbon emission for an EV. It is worth noting that scaling up the value of E_0 is equivalent to scaling down the value of ϕ . For example, the ride assignment of TORA with $\phi = 10$ and $E_0 = 63.35g.CO_2eq/km$ equals to the ride assignment of TORA with $\phi = 5$ and $E_0 = 126.7g.CO_2eq/km$.

⁴We used the Tesla Model Y as a prototype EV, which has an energy efficiency rating of 26 kWh/100mi. We used the average carbon intensity value for Austin, Texas, which is 408 g.CO₂eq/kWh to compute the unit distance emissions.

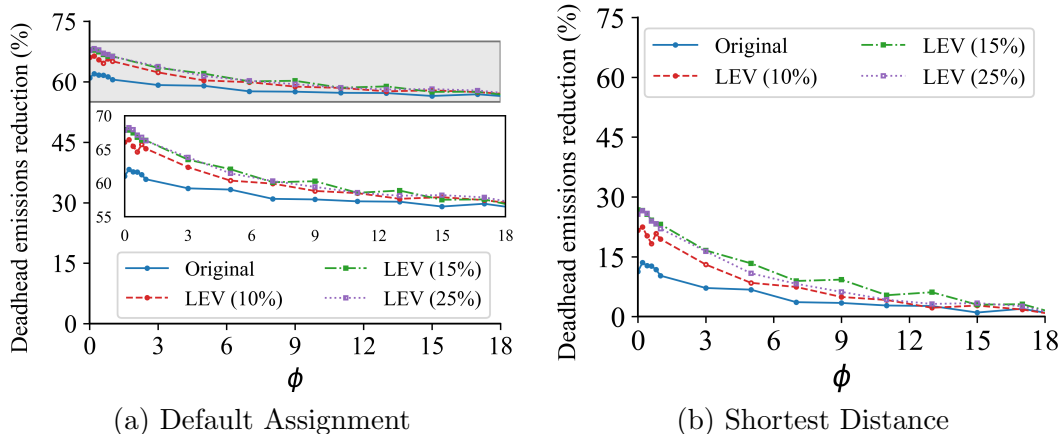


Figure 5: Percentage reduction in deadhead miles emissions on y -axis compared to the (a) default assignment in the RideAustin dataset and (b) shortest distance driver assignment, as a function of ϕ on the x -axis for different LEV fractions. Lower values of ϕ generally yield higher reduction.

Performance and equity metrics. We assess ride assignment strategies using three performance metrics: passenger waiting time, deadhead emissions, and overall ridesharing emissions.

To analyze the equity implications of the ride assignments algorithms, we leverage the parameter ϕ . Besides the previous LEVs, we introduce another category of high-emission vehicles (HEVs) with lower than 11.7 liters per 100km (20mpg). We report the fraction of trips assigned and the average deadhead-to-trip distance ratio for the two categories. For an equitable outcome, the values for both metrics should be similar across vehicle categories.

6.2 Impact on Emissions and Waiting Times

In this section, we present the outcomes of our experiments, which shed light on the intricate interplay of ride assignment algorithms, routing strategies, and various dataset configurations. These results offer valuable insights into the performance trade-offs inherent in ridesharing platforms. We report our results by varying the threshold values (ϕ) for different datasets, including the original vehicle dataset and synthetically generated datasets with the injection of 10%, 15%, and 25% low-emission vehicles (LEVs). The objective is to investigate the impact of threshold parameter ϕ on two pivotal factors: the reductions in deadhead emissions (in Figure 5) and the increase in waiting times of the riders (in Figure 6).

The results in Figure 5 show that smaller values of threshold value ϕ lead to greater improvements in deadhead emissions. Additionally, at a fixed threshold value, a higher percentage of low-emission vehicles results in comparatively greater improvements in deadhead emissions. Regarding waiting times, as shown in Figure 6, a smaller threshold value increases the relative waiting time for both the default assignment and the shortest distance assignment strategies. Also, as ϕ increases, the relative increase in waiting time decreases. We note that with higher values of ϕ , our results demonstrate a negative increase in waiting time, i.e., a reduction in waiting time, compared to the default ride assignment in the RideAustin dataset. This happens because the default assignment does not always assign the closest drive to the passenger, and there is an opportunity to reduce waiting time without increasing emissions.

Key takeaway 1. *There is a trade-off between reducing the average deadhead emission and reducing the average passenger’s waiting time under the ride assignment of TORA, e.g., while smaller threshold values substantially reduce the emissions, it comes at the expense of increased waiting time of riders.*

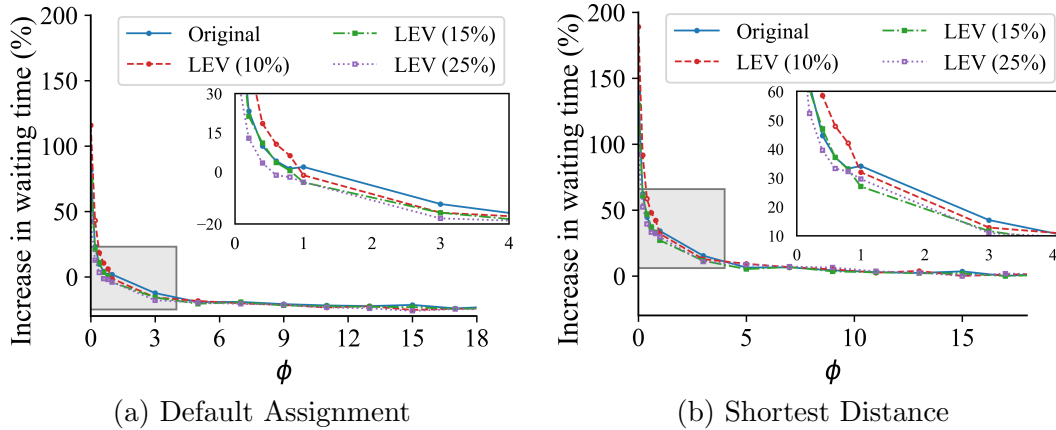


Figure 6: Percentage increase in waiting time on the y -axis compared to the (a) default assignment in the RideAustin dataset and (b) shortest distance driver assignment, as a function of ϕ on the x -axis for different LEV fractions. Higher values of ϕ generally yield lower waiting time.

6.3 Equity Implications

Depending on how we look at the problem and who we ask, the definition of equity can significantly change. In our context, we define it as an equitable distribution of rides across all vehicles. Since emission-aware ride assignment strategies favor high-efficiency vehicles, they can starve drivers with low efficiency, which typically belong to low-income drivers. As a result, there is a tradeoff between reduction in emissions and equity. In this section, we navigate the tradeoff using various values of ϕ . Note that any other definition of equity, such as the waiting time experienced by riders from low-income communities vs. affluent neighborhoods, can be explored using the information augmented into the dataset by E²-RideKit.

As illustrated in Figure 7, adjusting the threshold value results in a trade-off: an increase in the threshold leads to a decrease in the percentage of trips assigned to LEVs, while the percentage of rides assigned to HEVs increases. This implies that a higher threshold favors HEV rides, a trend consistent across datasets with varying LEV percentages. Intriguingly, our results also demonstrate that a higher percentage of LEVs within a dataset corresponds to a greater percentage of trips completed by LEVs, irrespective of the threshold value. Notably, the original dataset, across a range of threshold values from 0.001 to 18, assigned between 7.6% to 16.6% of rides to LEVs. These percentages shifted to 20.7% to 66.8%, 25.5% to 56.2%, and 39.4% to 65.8% for datasets containing 10%, 15%, and 25% LEVs, respectively. This underscores the pivotal role of a higher proportion of LEVs in achieving environmentally friendly ridesharing outcomes within defined criteria.

A distinct facet of our investigation involved analyzing the ratio between deadhead and total trip distance, known as the deadhead-to-trip ratio. Figure 8 showcases the results as the threshold values vary for LEVs and HEVs. As expected, an increase in the threshold value is associated with a decrease in the deadhead-to-trip ratio, consistent across scenarios with a fixed percentage of LEVs.

Interestingly, our results demonstrate that raising the threshold for a fixed percentage of LEVs is linked to an increased deadhead-to-trip ratio for HEVs. Additionally, for a fixed threshold value, a higher percentage of LEVs results in a lower deadhead-to-trip ratio, reflecting improved efficiency. These observations emphasize the intricate relationship between ride assignment parameters and the efficiency of trip assignments in diverse environmental contexts. Results show that the deadhead-to-trip ratio for LEVs varies between 31.8% to 35.7% for original vehicles when threshold varies between 0.001 to 18 while these ranges for datasets with 10%, 15%, and 25% LEVs were 32.8% to 37.0%, 32.4% to 34.0%, and 31.9% to 34% respectively.

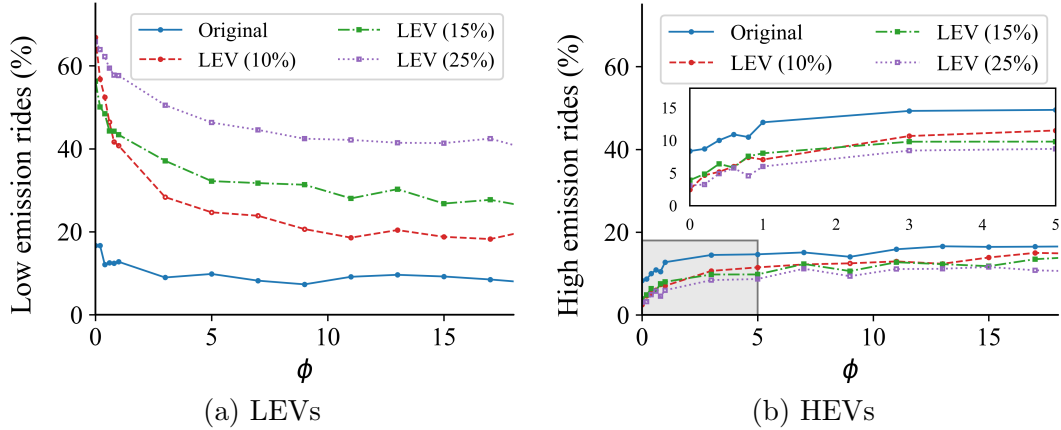


Figure 7: Fraction of rides assigned on y -axis to (a) LEVs and (b) HEVs, as a function of ϕ on the x -axis for different LEV fractions. At lower ϕ values, TORA assigns more rides to LEVs.

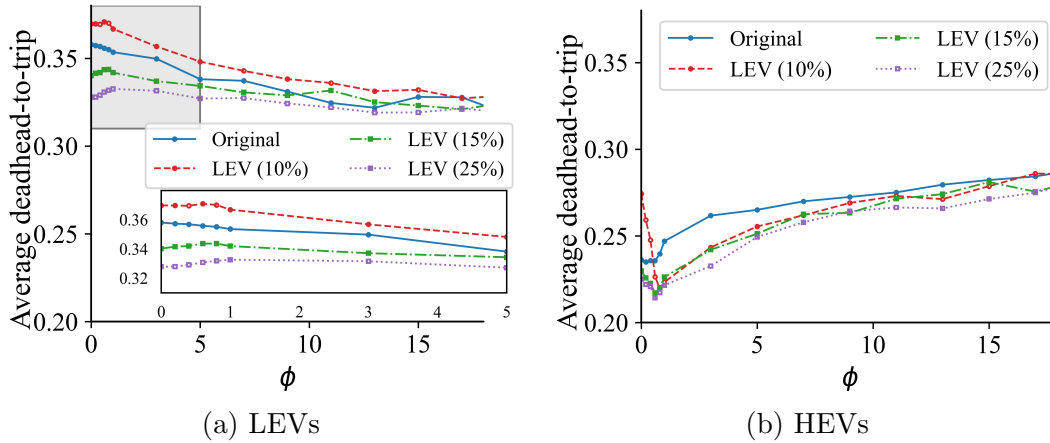


Figure 8: Average deadhead-to-trip distance ratio on y -axis for (a) LEVs and (b) HEVs, as a function of ϕ on the x -axis for different LEV fractions. At higher ϕ values, the deadhead-to-distance ratio for LEVs decreases but increases for HEVs.

Key Takeaway 2. Under the ride assignment of TORA, reduction in deadhead emissions requires sacrificing equity among drivers. In this case, a greater fraction of passengers get assigned to LEVs, and their deadhead miles would be longer.

6.4 Ride Assignment and Routing Comparison

Finally, we examine the impact of ride assignments and routing options on total emissions and waiting times, considering datasets with varying percentages of LEVs. By leveraging data from the Google Maps API, including travel times for different route options, the shortest (S), the fastest (F), and the fuel-efficient (E), explained in Section 4, we evaluated the performance of TORA with three different values of ϕ : 0.1, 1, and 7.5. Our results, illustrated in Figure 9, indicate when we increase the threshold value, we observe a decrease in the percentage increase in waiting time. Higher thresholds are associated with a reduced percentage increase in waiting times and a corresponding decrease in improvements in deadhead emissions. Conversely, lowering the threshold value amplifies waiting time increases and enhances the reduction of deadhead emissions. These findings illuminate the intricate dynamics between ride assignment and routing strategies and their consequential effects on emissions and waiting times within ridesharing systems. In addition, results reveal that the impact of ride assignment (different threshold values) on the total emission and waiting times is substantially greater than that of routing algorithms. Note with original vehicles, changing the threshold value from 7.5

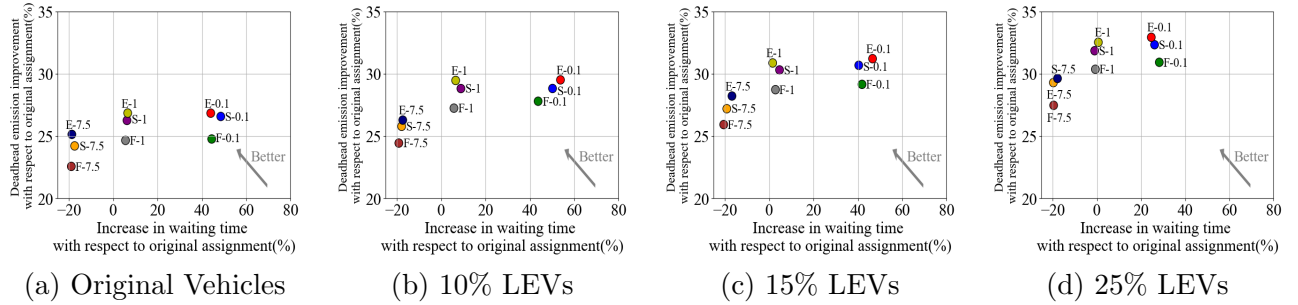


Figure 9: Percentage of improvement in total emission as a function of percentage of increase in waiting time assignment for TORA using three different route options (Shortest (S), Fastest (F), and Fuel Efficient (E)) and varying ϕ values: 0.1, 1, and 7.5.

to 0.1 can improve the total emission up to 3.5%, and increase the waiting time up to 71.2%. In contrast, for a fixed ride assignment algorithm, the maximum impact of the routing algorithm on the improvement of emission and increase in waiting time is 2.6%, and 10.1%.

Key Takeaway 3. *The impact of ride assignment algorithms on the emission and waiting time of ride sharing service is significantly more than the impact of routing algorithms.*

7 Related Work

Prior work on improving ride-sharing services has mainly focused on investigating the factors affecting passenger demand and scaling drivers’ availability to match demand by leveraging either theoretical approaches or empirical and ML-based approaches.

Theoretical approaches. Feng et al. [33] present a novel two-stage stochastic matching model for ride-hailing platforms, addressing uncertainty in rider and driver availability. Li et al. [34] study the strategies taken by taxi drivers for finding passengers from a large-scale taxi GPS dataset. Vazifeh et al. [35] propose an optimization-based approach to minimize the fleet size in on-demand urban mobility services. Zha et al. [36] analyze the ride-sourcing market by leveraging a model that captures the matchings between customers and drivers through a matching function. Bai et al. [37] and Feng et al. [38] propose a model for an on-demand service platform to estimate the passenger’s queuing time or their matching time. Abkarian et al. [39] present a model that aims to balance the trade-off between waiting times for reservation-based and on-demand users while minimizing the overall deadhead mileage driven by the vehicles in the fleet. Other studies analyze ride-sharing for deadheading, cost, energy consumption, and environmental impact through modeling, simulations, and data analytics [40, 41, 42].

Empirical and ML-based approaches. Moreira-Matias et al. [43] devise a method to predict the short-term distribution of taxi passengers using streaming data from taxis operating in Porto, Portugal. Iacobucci et al. [44] investigate the potential demand for Shared Autonomous Vehicles (SAVs) or robotaxis using a scalable simulation framework. This study explores the impact of SAVs on travel behavior, considering factors such as fare, waiting times, and real-time demand. Jungel et al. [45] focus on the development of online control algorithms for autonomous mobility-on-demand systems. They propose a hybrid combinatorial optimization enriched machine learning pipeline that learns online dispatching and rebalancing policies from optimal full-information solutions.

Lavieri et al. [46] use an Austin-based ride-sharing dataset to present two statistical models: (i) a spatially lagged multivariate count model to estimate the number of trips generated in a specific zone during weekdays and weekends and (ii) fractional split model to identify the key features of zones where a majority of ride-sharing trips start. In another study, Liu et al. [47] analyze the temporal and spatial

patterns of riders’ requests and the ride services they choose. They use Random Forests to predict the travel demand of these ride-sharing services. Ke et al. [48] propose a novel spatio-temporal deep learning approach that uses a convolutional neural network (CNN) to model the spatial distribution of demand and a long short-term memory (LSTM) network to model the temporal patterns in ride demand. Other studies have also used LSTM networks to predict demand behaviors across both spatial and temporal dimensions [49, 50, 51]. While these studies focus on improving the performance of ride-sharing services, they do not explicitly target reducing deadhead miles. To the best of our knowledge, the most relevant work to our proposal is done by Kontou et al. [23]. Authors show up to 82% reduction in *trip-level* deadhead miles by leveraging hour-ahead trip demand predictions and a heuristic approach to driver assignment. However, their focus on reducing trip-level deadhead miles may not always lead to reducing system-wide deadhead miles and emissions that, in addition to the number of miles, depend on the fuel efficiency of vehicles and traffic conditions. Furthermore, they do not consider equity metrics from the rider’s or driver’s perspective.

Our paper is the first to take a holistic approach toward designing emission-aware ride assignment optimizations that explicitly target reducing emissions from deadhead miles, embed equity considerations into the ride assignment process, and consider electric and low-emission vehicles as a part of ride-sharing fleets. In addition, our solution approach is a data-driven algorithmic design that relies on a straightforward greedy algorithm.

Lastly, we note that there is extensive literature on optimizing carbon emissions and energy consumption for vehicle routing. The existing emission-optimized routing algorithms, e.g., [52, 53, 54], mainly model the problem as extended variants of the classic traveling salesman problem under different settings, e.g., joint electric vehicle charging and routing [55, 56, 57, 58], joint routing and speed optimization [59, 52], and environmental impacts of the routing decisions [60, 61]. In this work, however, we used off-the-shelf routing algorithms from Google Map API. An interesting future direction of our work is to combine the above routing algorithms with the ride assignment policies presented in this paper.

8 Concluding Remarks

In this paper, we presented equity- and emission-aware ride assignment and routing approaches to reduce the overall emissions of the ridesharing platform. We then presented **E²-RideKit** as a toolkit that combines multiple datasets needed for the performance evaluation of our algorithms. Future research should focus on explicitly integrating equity considerations into the ride assignment optimization problem. This involves developing models that account for various dimensions of equity, such as fairness in service access and drivers’ treatment, to align the optimization process with broader social goals. Another critical area for further exploration is a thorough analysis of the simple online algorithm used in this study. Investigating its robustness, scalability, and sensitivity to different conditions will provide valuable insights, enabling potential refinements and improvements to enhance its performance under various scenarios. Additionally, beyond carbon minimization, future research should extend the evaluation framework to include a comprehensive assessment of the impact on drivers’ income. Balancing environmental objectives with socioeconomic considerations is essential to ensure that emission reduction efforts do not compromise the livelihoods of those providing ride-sharing services. Finally, we will publicly release **E²-RideKit** as an open-source toolkit to facilitate research in emissions analysis of ridesharing ecosystems.

9 Acknowledgements

This research was supported in part by NSF grants CAREER 2045641, CNS-2102963, CNS-2106299, CPS-2136199, and NGSDI-2105494.

References

- [1] Congressional Budget Office. Emissions of carbon dioxide in the transportation sector. <https://www.cbo.gov/publication/58566>, 2022.
- [2] International Energy Agency. World energy outlook 2022. <https://www.iea.org/reports/world-energy-outlook-2022>, 2022.
- [3] Directorate-General for Mobility and Transport, European Commission. Urban mobility. https://transport.ec.europa.eu/transport-themes/clean-transport-urban-transport/urban-mobility_en, Accessed March 2023.
- [4] Jimi B. Oke, Youssef M. Aboutaleb, Arun Akkinepally, Carlos Lima Azevedo, Yafei Han, P. Christopher Zegras, Joseph Ferreira, and Moshe E. Ben-Akiva. A novel global urban typology framework for sustainable mobility futures. *Environmental Research Letters*, 14(9):095006, 2019.
- [5] Ride sharing market size to hit usd 205.83 billion, globally, by the end of 2030, growing at a cagr of 13.5%. <https://tinyurl.com/BloombergRideSharing>, 2023. Access: 2023-05-20.
- [6] Shan Jiang, Le Chen, Alan Mislove, and Christo Wilson. On ridesharing competition and accessibility: Evidence from Uber, Lyft, and taxi. In *Proceedings of the 2018 World Wide Web Conference*, pages 863–872, 2018.
- [7] Knowledge at Wharton Staff. How green is the sharing economy? <https://knowledge.wharton.upenn.edu/article/how-green-is-the-sharing-economy/>, 2020. (Accessed on 05/05/2020).
- [8] Alejandro Henao. *Impacts of Ridesourcing-Lyft and Uber-on Transportation Including VMT, Mode Replacement, Parking, and Travel Behavior*. University of Colorado at Denver, 2017.
- [9] Laura Bliss. How much traffic do Uber and Lyft cause? <https://www.citylab.com/transportation/2019/08/uber-lyft-traffic-congestion-ride-hailing-cities-drivers-vmt/595393/>, August 2019.
- [10] Bruce Schaller. Unsustainable? The growth of app-based ride services and traffic, travel and the future of New York City. Technical report, Schaller Consulting, New York, NY, December 2017.
- [11] Don Anair, Jeremy Martin, Maria Cecilia Pinto de Moura, and Joshua Goldman. Ride-hailing’s climate risks. <https://www.ucsusa.org/resources/ride-hailing-climate-risks>, 2020. (Accessed on 05/05/2020).
- [12] Ride-hailing increases emissions, contributes to climate pollution - study. <https://www.reuters.com/article/us-uber-emissions/ride-hailing-increases-emissions-contributes-to-climate-pollution-study-idUSKBN20J27K>, 2020. (Accessed on 05/05/2020).
- [13] Anurag Komanduri, Zeina Wafa, Kimon Proussaloglou, and Simon Jacobs. Assessing the impact of app-based ride share systems in an urban context: Findings from Austin. *Transportation Research Record*, 2672(7):34–46, 2018.

- [14] Alejandro Henao and Wesley E Marshall. The impact of ride-hailing on vehicle miles traveled. *Transportation*, 46(6):2173–2194, 2019.
- [15] Judd Cramer and Alan B Krueger. Disruptive change in the taxi business: The case of uber. *American Economic Review*, 106(5):177–182, 2016.
- [16] Alex Rosenblat, Karen Levy, Solon Barocas, and Tim Hwang. Discriminating tastes: Customer ratings as vehicles for bias. *Data & Society*, pages 1–21, 2016.
- [17] Cody Cook, Rebecca Diamond, Jonathan V Hall, John A List, and Paul Oyer. The gender earnings gap in the gig economy: Evidence from over a million rideshare drivers. *The Review of Economic Studies*, 88(5):2210–2238, 2021.
- [18] Mark Graham and Joe Shaw. *Towards a fairer gig economy*. Meatspace Press, 2017.
- [19] Andrés Fielbaum. Optimizing a vehicle’s route in an on-demand ridesharing system in which users might walk. *Journal of Intelligent Transportation Systems*, 26(4):432–447, 2022.
- [20] Dujuan Wang, Qi Wang, Yunqiang Yin, and TCE Cheng. Optimization of ride-sharing with passenger transfer via deep reinforcement learning. *Transportation Research Part E: Logistics and Transportation Review*, 172:103080, 2023.
- [21] Abood Mourad, Jakob Puchinger, and Chengbin Chu. A survey of models and algorithms for optimizing shared mobility. *Transportation Research Part B: Methodological*, 123:323–346, 2019.
- [22] RideAustin. Ride-austin-june6-april13. <https://data.world/ride-austin/ride-austin-june-6-april-13>, 2017.
- [23] Eleftheria Kontou, Venu Garikapati, and Yi Hou. Reducing ridesourcing empty vehicle travel with future travel demand prediction. 121:102826, 2020.
- [24] Google. gmapapi. <https://developers.google.com/maps>, 2023. Accessed: 2023-09.
- [25] Electricity Map. <https://www.electricitymap.org/map>, Accessed September 2020.
- [26] How clean is the electricity I’m using right now? <https://www.wattpower.org/>, 2023. Accessed: 2023-02-27.
- [27] U.S. Census Bureau. Vehicle emissions forecasting. <https://www.kaggle.com/datasets/abhikdas2809/canadacaremissions>, 2023.
- [28] U.S. Census Bureau. Measuring America’s people, places, and economy. <https://www.census.gov/>, 2023.
- [29] Geocodio’s restful api. <https://www.geocodio.io/docs/>, Accessed September 2023.
- [30] United States Census Bureau. Census data, 2020.
- [31] Stuart Lloyd. Least Squares Quantization in PCM. *IEEE Transactions on Information Theory*, 1982.
- [32] U.S. Department of Energy. Electric vehicle charging station locations. https://afdc.energy.gov/fuels/electricity_locations.html#/find/nearest?fuel=ELEC, 2023.
- [33] Yiding Feng, Rad Niazadeh, and Amin Saberi. Two-stage stochastic matching with application to ride hailing. In *Proceedings of the 2021 ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 2862–2877. SIAM, 2021.

- [34] Bin Li, Daqing Zhang, Lin Sun, Chao Chen, Shijian Li, Guande Qi, and Qiang Yang. Hunting or waiting? discovering passenger-finding strategies from a large-scale real-world taxi dataset. In *2011 IEEE International Conference on Pervasive Computing and Communications Workshops (PERCOM Workshops)*, pages 63–68. IEEE, 2011.
- [35] Mohammad M Vazifeh, Paolo Santi, Giovanni Resta, Steven H Strogatz, and Carlo Ratti. Addressing the minimum fleet problem in on-demand urban mobility. *Nature*, 557(7706):534–538, 2018.
- [36] Liteng Zha, Yafeng Yin, and Hai Yang. Economic analysis of ride-sourcing markets. *Transportation Research Part C: Emerging Technologies*, 71:249–266, 2016.
- [37] Jiaru Bai, Kut C So, Christopher S Tang, Xiqun Chen, and Hai Wang. Coordinating supply and demand on an on-demand service platform with impatient customers. *Manufacturing & Service Operations Management*, 21(3):556–570, 2019.
- [38] Guiyun Feng, Guangwen Kong, and Zizhuo Wang. We are on the way: Analysis of on-demand ride-hailing systems. *Manufacturing & Service Operations Management*, 23(5):1237–1256, 2021.
- [39] Hoseb Abkarian, Hani S Mahmassani, and Michael Hyland. Modeling the mixed-service fleet problem of shared-use autonomous mobility systems for on-demand ridesourcing and carsharing with reservations. *Transportation Research Record*, 2676(8):363–375, 2022.
- [40] Gordon S Bauer, Jeffery B Greenblatt, and Brian F Gerke. Cost, energy, and environmental impact of automated electric taxi fleets in Manhattan. *Environmental science & technology*, 52(8):4920–4928, 2018.
- [41] Gopindra S Nair, Chandra R Bhat, Irfan Batur, Ram M Pendyala, and William HK Lam. A model of deadheading trips and pick-up locations for ride-hailing service vehicles. *Transportation Research Part A: Policy and Practice*, 135:289–308, 2020.
- [42] Tom Wenzel, Clement Rames, Eleftheria Kontou, and Alejandro Henao. Travel and energy implications of ridesourcing service in Austin, Texas. *Transportation Research Part D: Transport and Environment*, 70:18–34, 2019.
- [43] Luis Moreira-Matias, Joao Gama, Michel Ferreira, Joao Mendes-Moreira, and Luis Damas. Predicting taxi-passenger demand using streaming data. *IEEE Transactions on Intelligent Transportation Systems*, 14(3):1393–1402, 2013.
- [44] Riccardo Iacobucci, Jonas Donhauser, Jan-Dirk Schmöcker, and Marco Pruckner. The demand potential of shared autonomous vehicles: a large-scale simulation using mobility survey data. *Journal of Intelligent Transportation Systems*, pages 1–22, 2023.
- [45] Kai Jungel, Axel Parmentier, Maximilian Schiffer, and Thibaut Vidal. Learning-based online optimization for autonomous mobility-on-demand fleet control. *arXiv preprint arXiv:2302.03963*, 2023.
- [46] Patrícia S Lavieri, Felipe F Dias, Natalia Ruiz Juri, James Kuhr, and Chandra R Bhat. A model of ridesourcing demand generation and distribution. *Transportation Research Record*, 2672(46):31–40, 2018.
- [47] Jiaokun Liu, Erjia Cui, Haoqiang Hu, Xiaowei Chen, Xiqun Chen, and Feng Chen. Short-term forecasting of emerging on-demand ride services. In *2017 4th International Conference on Transportation Information and Safety (ICTIS)*, pages 489–495. IEEE, 2017.

- [48] Jintao Ke, Hongyu Zheng, Hai Yang, and Xiqun Michael Chen. Short-term forecasting of passenger demand under on-demand ride services: A spatio-temporal deep learning approach. *Transportation research part C: Emerging technologies*, 85:591–608, 2017.
- [49] Jun Xu, Rouhollah Rahmatizadeh, Ladislau Bölöni, and Damla Turgut. Real-time prediction of taxi demand using recurrent neural networks. *IEEE Transactions on Intelligent Transportation Systems*, 19(8):2572–2581, 2017.
- [50] Donggyun Ku, Sungyong Na, Jooyoung Kim, and Seungjae Lee. Real-time taxi demand prediction using recurrent neural network. In *Proceedings of the Institution of Civil Engineers-Municipal Engineer*, volume 174, pages 75–87. Thomas Telford Ltd, 2021.
- [51] Chao Wang, Yi Hou, and Matthew Barth. Data-driven multi-step demand prediction for ride-hailing services using convolutional neural network. In *Advances in Computer Vision: Proceedings of the 2019 Computer Vision Conference (CVC), Volume 2 1*, pages 11–22. Springer, 2020.
- [52] Lei Deng, Mohammad H. Hajjesmaili, Minghua Chen, and Haibo Zeng. Energy-efficient timely transportation of long-haul heavy-duty trucks. In *Proc. of ACM International Conference on Future Energy Systems (eEnergy)*, 2016.
- [53] Qingyu Liu, Haibo Zeng, and Minghua Chen. Energy-efficient timely truck transportation for geographically-dispersed tasks. In *Proceedings of the Ninth International Conference on Future Energy Systems*, pages 324–339, 2018.
- [54] Wenjie Xu, Qingyu Liu, Minghua Chen, and Haibo Zeng. Ride the tide of traffic conditions: Opportunistic driving improves energy efficiency of timely truck transportation. *IEEE Transactions on Intelligent Transportation Systems*, 2023.
- [55] Guy Desaulniers, Fausto Errico, Stefan Irnich, and Michael Schneider. Exact algorithms for electric vehicle-routing problems with time windows. *Operations Research*, 64(6):1388–1405, 2016.
- [56] Dominik Goeke and Michael Schneider. Routing a mixed fleet of electric and conventional vehicles. *European Journal of Operational Research*, 245(1):81–99, 2015.
- [57] Alejandro Montoya, Christelle Guéret, Jorge E Mendoza, and Juan G Villegas. The electric vehicle routing problem with nonlinear charging function. *Transportation Research Part B: Methodological*, 103:87–110, 2017.
- [58] Michael Schneider, Andreas Stenger, and Dominik Goeke. The electric vehicle-routing problem with time windows and recharging stations. *Transportation science*, 48(4):500–520, 2014.
- [59] Junyan Su, Qiulin Lin, and Minghua Chen. Follow the sun and go with the wind: Carbon footprint optimized timely e-truck transportation. In *Proceedings of the 14th ACM International Conference on Future Energy Systems*, pages 159–171, 2023.
- [60] Sevgi Erdoğan and Elise Miller-Hooks. A green vehicle routing problem. *Transportation research part E: logistics and transportation review*, 48(1):100–114, 2012.
- [61] Tolga Bektaş and Gilbert Laporte. The pollution-routing problem. *Transportation Research Part B: Methodological*, 45(8):1232–1250, 2011.