RIDERS: Real-time Information Dissemination for Efficiency in a Robo-taxi System

August 2025

Region 1:
New England University Transportation Center

Name of PI: Monika Filipovska, Ph.D.
University of Connecticut

Other Researcher: Haimanti Bala
University of Connecticut

In cooperation with U.S. Department of Transportation,
Office of the Assistant Secretary for Research and Technology (OST-R)

Grant #: 69A3552348301

TECHNICAL DOCUMENTATION

1. Project No.	2. Government Accession No.	3. Recipient's Catalog No.
161137	01904459	
4. Title and Subtitle RIDERS: Real-time Information Dissemination for Efficiency in a Robo-taxi System		**5. Report Date** August 2025 **6. Performing Organization Code** N/A
7. Author(s) **Monika Filipovska**, Ph.D. ORCiD: 0000-0002-2718-4722; **Haimanti Bala** ORCiD: 0000-0001-7241-522X.		**8. Performing Organization Report No.** N/A
9. Performing Organization Name and Address New England University Transportation Center 181 Presidents Drive University of Massachusetts - Amherst Amherst, MA 01003		**10. Work Unit No. (TRAIS)** **11. Contract or Grant No.** 69A3552348301
12. Sponsoring Agency Name and Address United States Department of Transportation Research and Innovative Technology Administration 1200 New Jersey Avenue, SE Washington, DC 20590		**13. Type of Report and Period Covered** Final Research Report **14. Sponsoring Agency Code** USDOT OST-R

15. Supplementary Notes

Report uploaded and accessible at the NEUTC Website (www.umass.edu/neutc)

16. Abstract

Real-time traffic information is critical for efficient robo-taxi fleet management, yet such data is often costly to obtain and incomplete in mixed traffic with low autonomous vehicle (AV) penetration. This project addresses this challenge by proposing a hierarchical reinforcement learning (HRL) framework that enables fleets to act not only as mobility providers but also as mobile sensors. Using New York City taxi demand data on the Manhattan road network, we simulate fleets that begin with partial knowledge of link-level speeds. Two actor–critic agents are employed: a zone-level agent reallocates idle vehicles across demand zones, while a route-level agent selects paths that balance service efficiency with information gain. Numerical experiments compare an information-focused strategy (AVR-IF) against a baseline (AVR-BA) under weekday and weekend demand with varying fleet sizes. Results show that AVR-IF uncovers vastly more link-level traffic information (~75,000 observations vs. ~3,000 for the baseline) while maintaining similar passenger wait times and vehicle miles traveled. These findings demonstrate that fleets can generate valuable real-time network intelligence without increasing operational costs or user burden. Beyond improving adaptability and responsiveness, such dual-purpose operation reduces reliance on third-party data and creates opportunities for collaboration between operators and public agencies. Ultimately, integrating information collection into everyday operations can make robo-taxi fleets more effective and impactful than traditional taxi or ride-sourcing services, advancing both mobility efficiency and urban traffic management.

17. Key Words	18. Distribution Statement
Robo-taxi, Mobility-on-Demand, Operations, Information Collection, Connected and Automated Vehicles	No restrictions.

19. Security Classif. (of this report)	20. Security Classif. (of this page)	21. No. of Pages	22. Price
Unclassified	Unclassified	4	

Form DOT F 1700.7 (8-72) **Reproduction of completed page authorized.**

About NEUTC

The New England Regional University Transportation Center (NEUTC) is a diverse, multidisciplinary consortium committed to addressing the pressing issue of traffic safety. Our objective, in line with the Infrastructure Investment and Jobs Act (IIJA), is to drive transformative research, education, and technology transfer to address critical traffic safety needs in a time when roadway fatalities are distressingly high.

Our research and educational activities at NEUTC are guided by four principal safety themes, each addressing a critical challenge in transportation safety. These themes capture the various integral components of the transportation system, focusing on technology, infrastructure, vehicles, and users with a commitment to safety and public engagement. Our overarching theme is promoting safety, with the common underlying science being the study of behavioral, systemic, environmental, and mobility-driven factors on safety.

Disclaimer

The contents of this report reflect the views of the authors, who are responsible for the facts and the accuracy of the information presented herein. This document is disseminated in the interest of information exchange. The report is funded, partially or entirely, under 69A3552348301 from the U.S. Department of Transportation's University Transportation Centers Program. The U.S. Government assumes no liability for the contents or use thereof.

All matching funds were University salaries and graduate tuition waivers.

Acknowledgments

This work is supported by the New England University Transportation Center (NEUTC). The authors remain fully responsible for all contents of the report, which may not necessarily reflect the views of the sponsoring agencies.

Motivation

Autonomous vehicles (AVs) are poised to transform mobility through connectivity, electrification, shared use, and on-demand services. Connected and automated vehicles (CAVs), in particular, enable wireless communication with vehicles, infrastructure, and devices (Ahmed et al. 2022; Atkins 2016; Wang et al. 2021). This capability becomes particularly valuable in the context of Robo-taxis, a system in which CAVs operate as part of a shared, often on-demand, mobility platform designed to serve multiple users efficiently, either concurrently (e.g. ride-splitting) or sequentially (e.g. ride-hailing) (Bala et al. 2023).

In fleet operations, access to real-time traffic data is critical for routing, assignment, and repositioning. However, traffic information is often costly and incomplete, especially in mixed traffic where CAV penetration is low. This limits continuous data collection and reduces reliability of information exchange. A promising solution is to strategically route idle robo-taxis through areas with limited coverage, allowing the fleet itself to generate and update traffic data to support more efficient, data-driven decisions.

Research on leveraging CAV capabilities for route guidance remains limited. Earlier studies have explored real-time route guidance, traveler decision-making in mixed traffic, and safety improvements through vehicle-to-vehicle communication (Genders and Razavi 2016; Tian et al. 2013; Yang and Liu 2022). The most relevant study was done by Liu and coauthors, who examined information use in a shared autonomous taxi system, combining historical and real-time traffic data for pathfinding, though their system focused only on assignment and service without addressing repositioning (Liu et al. 2018).

This study intends to provide a distinctive routing strategy for robo-taxi fleet that prioritizes collecting the most recent information from road segments or links, leveraging its connectivity to actively collect and update traffic information across the service region. A hierarchical reinforcement learning (HRL) framework is employed with two agents: a zone selection agent that allocates idle vehicles to different zones within the service region, and a route selection agent that aims to improve data collection coverage. By actively collecting and updating traffic information, the system enhances coverage, reduces travel times, and improves demand responsiveness, thereby advancing the efficiency and robustness of robo-taxi fleet operations.

Executive Summary

Methodology

The study models a centralized robo-taxi fleet management system, where autonomous vehicles (AVs) are assigned to serve dynamically arriving trip requests within an urban service area divided into zones. At any given time, both vehicles and requests belong to one zone. Operations focus on two key functions:

- Passenger-vehicle assignment: handled at regular intervals using a First-Come-First-Serve with Smart Nearest Neighbor (FCFS-smartNN) policy, which matches idle vehicles to passengers by request order and proximity. Each vehicle serves one passenger per trip (Hyland and Mahmassani 2018).

- Vehicle repositioning: idle vehicles are periodically relocated between zones to balance fleet distribution. They travel directly to the centroid of their assigned zone and are immediately available for new requests upon arrival.

A hierarchical reinforcement learning (HRL) framework is adopted to manage large-scale robo-taxi fleet repositioning. The framework is designed to capture both high-level fleet balancing decisions and lower-level routing choices that together influence service quality and information coverage. The decision-making is organized into two cooperative layers:

- Zone Selection Agent (high-level): This agent reallocates idle vehicles across zones to better match supply with demand. At regular intervals, it determines what fraction of idle vehicles in each zone should move to other zones, effectively rebalancing the fleet. The output is a set of repositioning flows between zones, ensuring that under-served areas receive more vehicles while avoiding oversupply elsewhere.
- Route Selection Agent (low-level): Once a repositioning decision is made, this agent determines which specific routes the vehicles should take. Rather than always selecting the shortest path, it considers the value of traversing under-observed links, thereby improving the system's knowledge of real-time traffic conditions. In this way, vehicles act not only as service providers but also as mobile sensors that continuously reduce gaps in network information.

The state space for decision-making includes idle vehicle distributions, zone-level demand, and link-level information availability. As vehicles traverse links, they collect and update traffic data, reducing information deficiencies and increasing the accuracy of the system's network knowledge.

The reward structure balances multiple objectives:

- At the zone level, rewards encourage supply-demand balance across regions by penalizing mismatches between demand and idle vehicle availability.
- At the route level, rewards penalize passenger wait times while also rewarding information gain from probing under-observed links.

The system learns policies that maximize long-term performance by applying an Advantage Actor–Critic (A2C) algorithm, with zone-level and route-level agents updated separately at different frequencies. Zone decisions are updated more frequently, while route decisions are aggregated and refined over longer training periods.

Through this hierarchical setup, the fleet learns to balance immediate service efficiency with the longer-term benefits of improved traffic knowledge, achieving a strategy that enhances both responsiveness and system robustness.

Numerical Experiments

The framework is evaluated using New York City yellow taxi trip data from April 2016 to represent ride demand and a simplified Manhattan road network (Wollenstein-Betech et al. 2021) for road topology. Traffic conditions are modeled with hourly link speeds. At the start of each simulation, vehicles only have partial traffic knowledge, which is updated as they traverse links and learn the true speeds.

Fleet operations are simulated with a custom agent-based tool that models passenger assignments, vehicle repositioning, routing, and state transitions. Pick-up and drop-off times are included, and simulations cover full operating days beginning at 7:00 a.m. Fleet management is controlled by the HRL framework with two A2C agents: a zone-level agent for inter-zone rebalancing and a route-level agent for path choices. Demand forecasts are refreshed every five minutes. To keep computation manageable, demand is scaled to 1% of historical requests, corresponding to different fleet sizes.

The experiments are designed to answer the following questions:

1. Does the proposed information-focused strategy (AVR-IF) improve efficiency compared to a baseline (AVR-BA) without the information-gathering component?
2. How robust is the proposed approach to changes in fleet size?
3. How sensitive is the strategy's performance to changes in demand patterns between weekdays and weekends?

Both models were tested with fleets of 60, 120, and 240 vehicles (with training conducted on the 60-vehicle case) and evaluated under both weekday and weekend demand patterns from NYC taxi data. This setup produced six scenarios in total, reflecting different combinations of fleet size and demand conditions, and capturing the temporal and spatial variations in travel behavior.

Results & Discussion

The comparison between the baseline model (AVR-BA) and the information-focused model (AVR-IF) highlights several important findings. First, both models visit a similar number of unique links in the Manhattan network (about 2,100–2,900, depending on fleet size) and achieve comparable coverage fractions, ranging from 67% to 93%. However, AVR-IF uncovers far more traffic information, around 75,000 link-level observations—while the baseline only discovers about 2,000–4,000. This confirms that prioritizing information collection drives vehicles to probe and update traffic conditions much more effectively, whereas the baseline remains narrowly demand-focused.

Second, this added exploration does not harm service quality. On weekdays, average daily passenger pickup wait times decrease as fleet size increases—from about 27 minutes with 60 vehicles to 23 minutes with 120 or more vehicles. In each case, AVR-IF performs slightly better than AVR-BA. By contrast, weekend operations show not much variation across fleet sizes or models. This indicates that weekend demand is lower and more evenly distributed, making repositioning less influential.

Finally, the analysis of vehicle miles traveled (VMT) shows that AVR-IF's information-seeking behavior does not inflate overall distance travelled by the fleet. As shown in Figure 1, weekday VMT is naturally higher than weekend VMT due to stronger demand, but differences between AVR-IF and AVR-BA remain negligible. For example, with a 120-vehicle fleet, weekday VMT averages about 7,480 miles for AVR-BA and 7,570 miles for AVR-IF, leading to a very small increase.

Overall, the results demonstrate that AVR-IF delivers substantial improvements in network knowledge while maintaining comparable passenger wait times and travel efficiency. The benefits are most visible under weekday conditions, when travel demand is higher and more uneven across the service region.

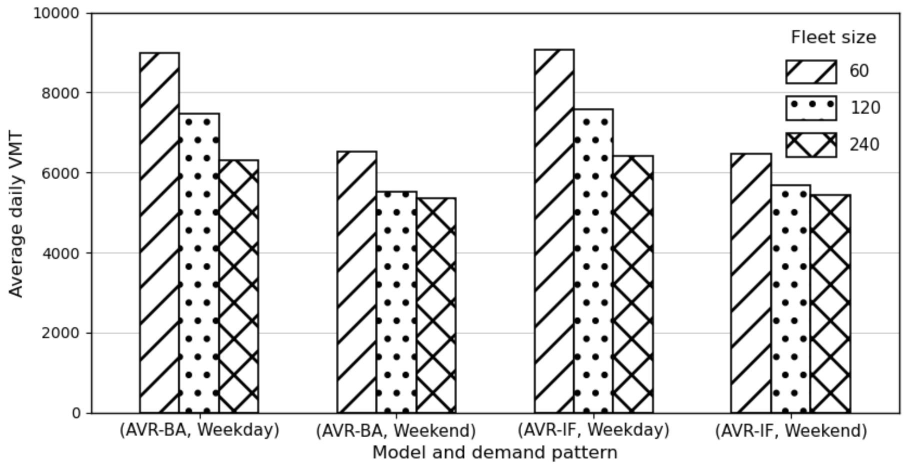

Figure 1. Bar plot of average daily VMT by models and demand patterns.

Outcomes

This study shows that robo-taxi fleets can integrate large-scale traffic information collection into daily operations without reducing service quality. The AVR-IF model consistently discovers far more link-level speeds than the baseline while keeping passenger wait times and VMT unchanged. This demonstrates that fleets can act as both mobility providers and mobile sensing systems at no extra cost to operators or inconvenience to users. Such a strategy reduces reliance on external traffic data, improves adaptability to dynamic conditions, and strengthens the case for shared autonomous fleets. Beyond operations, fleet-generated data can complement infrastructure-based sensing, offering cities new tools for congestion management, safety monitoring, and planning. Policymakers could build on this by promoting data-sharing frameworks to ensure that traffic intelligence benefits both operators and the public.

Impacts

The impacts span operational, passenger, system, environmental, and societal dimensions. For operators, integrating information collection reduces reliance on costly third-party data and supports more efficient deployment. Passengers benefit from dependable service and wait times while fleets simultaneously gather traffic data. At the system level, fleets as mobile sensors extend the reach of traffic intelligence without extra mileage, enabling smoother flow and more resilient networks. Environmentally, data are collected during normal trips, improving use of road space and reducing wasted energy. Societally, richer traffic data can enhance safety, guide smarter infrastructure investments, and strengthen planning capacity. Overall, operating robo-taxi fleets as both service providers and mobile sensors can deliver reliable rides and valuable network intelligence that benefits operators, passengers, and cities alike.

References

Ahmed, H. U., Y. Huang, P. Lu, and R. Bridgelall. 2022. “Technology Developments and Impacts of Connected and Autonomous Vehicles: An Overview.” *Smart Cities*, 5 (1): 382–404. https://doi.org/10.3390/smartcities5010022.

Atkins, W. S. 2016. “Research on the Impacts of Connected and Autonomous Vehicles (CAVs) on Traffic Flow.” *Stage 2: Traffic Modelling and Analysis Technical Report*. Department for Transport London, UK.

Bala, H., S. Anowar, S. Chng, and L. Cheah. 2023. “Review of studies on public acceptability and acceptance of shared autonomous mobility services: Past, present and future.” *Transp Rev*, 1–27. Taylor & Francis.

Genders, W., and S. N. Razavi. 2016. “Impact of connected vehicle on work zone network safety through dynamic route guidance.” *Journal of Computing in Civil Engineering*, 30 (2): 04015020. American Society of Civil Engineers.

Hyland, M., and H. S. Mahmassani. 2018. “Dynamic autonomous vehicle fleet operations: Optimization-based strategies to assign AVs to immediate traveler demand requests.” *Transp Res Part C Emerg Technol*, 92. https://doi.org/10.1016/j.trc.2018.05.003.

Liu, Z., T. Miwa, W. Zeng, M. G. H. Bell, and T. Morikawa. 2018. “Shared autonomous taxi system and utilization of collected travel-time information.” *J Adv Transp*, 2018. Hindawi.

Tian, D., Y. Yuan, J. Zhou, Y. Wang, G. Lu, and H. Xia. 2013. “Real-time vehicle route guidance based on connected vehicles.” *2013 IEEE International Conference on Green Computing and Communications and IEEE Internet of Things and IEEE Cyber, Physical and Social Computing*, 1512–1517. IEEE.

Wang, J., S. Jiang, Y. Qiu, Y. Zhang, J. Ying, and Y. Du. 2021. “Traffic Signal Optimization under Connected-Vehicle Environment: An Overview.” *J Adv Transp*, 2021: 1–16. https://doi.org/10.1155/2021/3584569.

Wollenstein-Betech, S., M. Salazar, A. Houshmand, M. Pavone, I. C. Paschalidis, and C. G. Cassandras. 2021. “Routing and rebalancing intermodal autonomous mobility-on-demand systems in mixed traffic.” *IEEE Transactions on Intelligent Transportation Systems*, 23 (8): 12263–12275. IEEE.

Yang, Z., and Y. Liu. 2022. “Optimal routing policy for a mixed traffic flow of connected vehicles and regular vehicles with en-route information.” *Available at SSRN 3885540*.