AVATAR: Autonomy Aware Routing for On-demand Transit Applications

David Rogers¹, Samir Gupta¹, Jose Paolo Talusan¹, Ammar Bin Zulqarnain¹, Mirza Baig², Arti Ramesh², Natsu Takahashi³, Naoki Kojo³, and Abhishek Dubey¹

¹Vanderbilt University ²Nissan North America, Inc. ³Nissan Motor Co., Ltd.

Abstract-Autonomous vehicles (AVs) are becoming integral to on-demand micro transit, offering the potential for safer, efficient, and sustainable transportation. However, AV deployment faces several challenges, including the lack of suitable roadways, varying travel conditions. Traditional routers prioritize speed and not reliability, leading to unpredictable operations and complications in planning. To address these, we introduce AVATAR, an autonomy-aware routing framework that prioritizes dependable, low-variance routes. Our approach encodes multiple objectives including road speed, speed variability, zoning areas, pedestrian encounters, and operator preferred roadways into edge-level routing engines. Objective optimized routes are generated, then scored using a multi-criteria decision-making process. User-configurable preference profiles, allow operators to define a balance between reliability and speed. AVATAR is a datadriven framework that supports both real-time AV operations and offline analysis, enabling transit operators to assess and refine routing strategies. Our experiments using real-world data from Silicon Valley, California, and Yokohama, Japan show that our approach significantly improves AV reliability and performance and advances the sustainable and scalable integration of AVs into future transportation networks.

Index Terms—Autonomous Vehicles, Path Planning, Optimization

I. INTRODUCTION

Autonomous vehicles (AVs) are increasingly becoming a cornerstone of on-demand micro-transit systems, promising to revolutionize urban transportation by offering safer, more efficient, and sustainable travel options [1]. The integration of AVs into these systems is driven by their potential to reduce traffic accidents, optimize fuel consumption, and provide consistent service [2, 3]. Despite these promising benefits, the deployment of AVs faces significant challenges that hinder their widespread adoption.

One of the primary challenges is the limited number of roadways suitable for AV operations. Unlike human drivers, AVs require highly reliable and predictable environments to function optimally. The current road infrastructure, with its varying conditions and frequent changes, poses a substantial obstacle [4]. For instance, roadways with high variance in travel conditions, such as fluctuating traffic patterns, construction zones, and unpredictable pedestrian encounters, can severely impact AV performance [5]. As shown in fig. 1, these factors create bottlenecks where AVs face low speeds,



Fig. 1. AV Partner observed lanes showing areas of low speed, high construction severity and high pedestrian encounters. Speed data for these observed lanes can vary between different times of the day and days of the week, resulting in fluctuating patterns.

heavy construction, and frequent pedestrians, complicating route planning.

Traditional routing algorithms, which are designed primarily for human drivers, prioritize travel time and often neglect the reliability and consistency required for AV operations [6]. This focus on speed can lead to unpredictable and inefficient routes, complicating the planning and execution of AV-based transit services. For example, a route that is fastest under ideal conditions might become highly unreliable during peak hours or in the presence of construction activities [7]. This unpredictability undermines the potential benefits of AVs and poses a significant barrier to their effective deployment. Additionally, the data needed to convey the reliability of a roadway at an adequate resolution is often difficult or expensive to obtain. For example, gathering detailed traffic patterns and construction reports necessitates extensive data collection efforts and continuous monitoring, making it costprohibitive for many transit operators [8].

To address these challenges, we propose an autonomy-aware routing framework, called AVATAR, that is tailored to AVs, prioritizing dependable, low-variance roadways by considering factors like road speed, speed variability, construction zones, pedestrian encounters, school zones, operator-preferred routes, and legally restricted roadways. Considering the limitations of existing work, this paper aims to address gaps by developing an autonomy-aware routing framework.

The summary of contributions are as follows:

 We design a multi-criteria decision-making process for holistic route evaluation based on user-configurable pref-

TABLE I Description of Symbols

Symbol	Description
A	Start location
B	End Location
R_{AB}	Set of Routes from A to B
r	A single route $\in R_{AB}$
r^*	Optimal route
C	Set of criteria to evaluate a route
c	A single criteria $\in C$
w	Weights for a criteria
S	Summary statistic function
α,β	Coefficients
σ	Standard deviation
E	Evaluation function

erences, balancing reliability and speed. AVATAR prioritizes dependable routes. Thus, is robust to environmental dynamics.

- Our approach incorporates a bootstrapping methodology using real-time General Transit Feed Specification (RT-GTFS) data from transit buses, providing insight into initial roadway conditions [9].
- AVATAR supports both real-time AV operations and offline analysis, it allows continuous assessment and refinement of routing strategies, significantly improving AV reliability and performance.
- Finally, we validate our approach with real-world data from Nashville, TN, Silicon Valley, California, and Yokohama, Japan.

Organization: This paper is divided into the following sections: We formulate the problem statement in Section II. Section III reviews the relevant literature. Section IV outlines the data collection and processing steps. Section V explains the proposed approach used in this paper. Section VI sets up the experiments and presents the key findings. Finally, Section VII discusses the implications of this work and potential future research directions.

II. PROBLEM STATEMENT

The deployment of autonomous vehicles (AVs) in ondemand microtransit systems presents a unique set of challenges that must be addressed to realize their full potential. One of the most pressing issues is the inadequacy of current routing methods, which do not account for the specific needs and constraints of AV operations. Traditional routing algorithms, designed primarily for human drivers, prioritize speed and often overlook the reliability and consistency required for AVs. This oversight leads to unpredictable and inefficient routes, complicating the planning and execution of AV-based transit services.

The core problem lies in the inability of existing routing methods to provide reliable and predictable routes for AVs. Unlike human drivers, AVs require highly reliable environments to function optimally. The current road infrastructure, characterized by varying conditions and frequent changes, poses a significant obstacle. Routes that are optimal in terms of speed under ideal conditions can become highly unreliable during peak hours or in the presence of other dynamic urban factors, leading to inconsistent AV performance and undermining the potential benefits of AVs.

To address these issues, it is crucial to contextualize the various factors affecting AV operations in the planning process. This involves integrating multiple operational objectives, such as road speed, speed variability, construction zones, pedestrian encounters, school zones, and operator-preferred routes, into the routing framework.

In the context of autonomous vehicle (AV) routing for ondemand micro transit, the goal is to generate, evaluate, and choose the most suitable route between two points, while considering multiple operational objectives. This problem can be formulated as a multi-objective optimization problem where a route is generated and evaluated based on a set of criteria that reflect the specific needs and constraints of AV operations. Table I summarizes all symbols used in the problem statement.

Let A and B be the starting and ending points, respectively. The objective is to find an optimal route from A to B that minimizes a weighted sum of various operational factors. The set of all possible routes from A to B is represented as R_{AB} . We evaluate each route r based on multiple criteria, such as road speed, speed variability, construction zones, pedestrian encounters, school zones, and routes preferred by the operator. Let $C = \{c_1, c_2, \ldots, c_n\}$ be the set of criteria used to evaluate each route. We formulate a naive evaluation of a route $r \in$ R_{AB} in eq. (1), where $c_i(r)$ denotes the value of the *i*-th criterion for the route r and w_i represents the weight assigned to the *i*-th criterion.

$$E(r) = \sum_{i=1}^{n} w_i \cdot c_i(r) \tag{1}$$

However, this formulation does not account for multiple representations of a distribution of a criterion, such as free-flow speed and constrained speed. The revised evaluation in eq. (2) addresses this limitation by incorporating a summary statistic $S_i(c_i(r))$, coefficients α_i and β_i , and standard deviation denoted by $\sigma(c_i(r))$ to balance tendency and variability.

$$E(r) = \sum_{i=1}^{n} w_i \cdot (\alpha_i \cdot \mathcal{S}_i(c_i(r)) + \beta_i \cdot \sigma(c_i(r)))$$
(2)

As described in eq. (3), our goal is to find the optimal route r^* that minimizes the weighted sum, taking into account both the performance and the consistency of each criterion.

$$r^* = \operatorname*{argmin}_{r \in R_{AB}} E(r) \tag{3}$$

III. RELATED WORK

Routing is a critical issue in traffic estimation systems and Intelligent Transportation Systems (ITS) in general. The challenge lies in developing routing algorithms that can adapt to real-world conditions and provide reliable, efficient routes for autonomous vehicles (AVs).

Dataset	Features	Frequency	Туре	Description	
OpenStreetMan	OSM Node	N/A	Spatial	Data point that represents a location on the map	
Opensucennap	OSM Edge	N/A	Spatial	Road segment between two OSM nodes	
	Lane Speed	N/A	Spatio-Temporal	Observed lane speeds at the 25th, 50th, and 75th percentiles	
AV Observations	Construction Observation	N/A	Spatio-Temporal	Detected construction zones along a lane segment	
	Pedestrian Observation	N/A	Spatio-Temporal	Detected pedestrian activity along a lane segment	
INRIX	Traffic Speed	15 minutes	Spatio-temporal	Median speed observed on a specific road segment	
	Bus ID	15-60 seconds	Temporal	Unique identifier for each bus	
	Bus Location	15-60 seconds	Spatial	Current latitude and longitude of the bus	
RT-GTFS	Bus Speed	15-60 seconds	Spatio-Temporal	Current speed of the bus (if available)	
	Bus Bearing	15-60 seconds	Spatio-Temporal	Current bearing of the bus (if available)	

TABLE II Data Features and Sources

Hu et al. [10] proposed an optimal route planning system for logistics vehicles based on artificial intelligence. Their approach leverages machine learning algorithms to predict traffic conditions and optimize routes for logistics operations. The system was evaluated using real-world data, demonstrating its effectiveness and efficiency in practical applications. However, their focus was primarily on logistics vehicles, and the approach does not fully address the unique challenges of AV routing in urban environments. Additionally, they do not incorporate any decision makers for choosing optimal routes based on use case.

Taha and AbuAli [11] discussed various route planning considerations for autonomous vehicles, emphasizing the importance of real-time data integration and the need for robust algorithms that can handle dynamic traffic conditions. Their work highlights the potential of cooperative tasks, such as speed harmonization and fuel consumption reduction, through aerodynamic drag reduction. While their study provides valuable insights into AV routing, it does not offer a comprehensive solution that integrates multiple operational objectives and real-time data sources.

Minh et al. [12] proposed effective traffic routing mechanisms aimed at enhancing urban transportation capacity and safety. Their approach adapts routing criteria based on geographical distance or estimated travel time, depending on user demands. The proposed mechanisms were shown to improve robustness, particularly in developing countries with less developed traffic infrastructures. However, their work primarily focuses on general traffic routing and does not fully address the specific requirements of AV operations.

Our proposed solution addresses these gaps by developing an autonomy-aware routing framework that prioritizes dependable, low-variance roadways and incorporates multiple operational objectives, such as road speed, speed variability, construction zones, pedestrian encounters, and school zones. By leveraging data from various sources, including RT-GTFS, INRIX, and AV data streams, our framework provides a comprehensive and adaptable routing solution for AVs.

The key advantage of our approach lies in its flexibility and robustness. Our framework can easily adapt to meet the diverse set of challenges on real-world roadways and quantify them in the context of dynamic operational requirements. This flexibility ensures that the routing decisions remain reliable and efficient, ultimately improving the performance and scalability of AV operations in on-demand micro transit systems.

IV. DATA COLLECTION AND PROCESSING

The development of an effective autonomy-aware routing framework requires collecting and processing large quantities of data that accurately reflect real-world road conditions. This section outlines the data sources, collection methods, and processing steps used to generate the datasets required for this study. A detailed list of all the features that were used as part of the routing framework is shown in Table II.

A. Data Sources

We collected data from various sources for Nashville, TN; the Bay Area, CA; and Yokohama, Japan. Below is a summary:

OSM Edges: OpenStreetMap (OSM) edges represent segments of the road network and are used as the base map for routing; each edge depicts a road segment between two OSM nodes (waypoints).

AV Partner Observations: This data was collected by AVs operating in the Silicon Valley Area. It includes lanes denoted by start and end coordinates with no direct mapping to OSM edges. Each lane is assigned a 0.25, 0.5, and 0.75 percentile speed, as well as data from a perception model that contains observations of construction and pedestrians.

INRIX: Provides speed data for major roads at a 15 minute intervals. We collect this data for the Nashville area. As INRIX data, shown in fig. 2, is not available for all roadways, experiments using this data are limited to areas with high INRIX availability.

RT-GTFS Data: Provides real-time information on the locations, speeds, and other attributes of transit vehicles. An extension of the static GTFS format, which includes scheduled transit data such as routes, stops, and timetables. We collect public RT-GTFS data for the Bay Area and Greater Tokyo Area, recording pings for each active bus every 15–60 seconds. Since our focus is on autonomous vehicles and shuttles, which operate at lower speeds than regular cars and can utilize transit lanes, RT-GTFS data serves as a reasonable proxy for estimating AV travel speeds. Figure 3 shows the RT-GTFS pings mapped to the OSM edges, highlighting areas with high and low variance.



Fig. 2. INRIX roadways in Nashville, TN colored by average speed.



Fig. 3. RT-GTFS pings mapped to OSM edges (blue) and annotated to show high variance and low variance edges.

B. Data Processing

The raw data requires several processing steps to ensure its accuracy and usability for routing purposes. The following methods are employed:

Filtering Bus Stops: We merge RT-GTFS data with static GTFS data in order to detect buses that are stopped or have low speed near scheduled bus stops. Where detected, we remove this data as it is not representative of general travel speeds in that location.

Calculating Missing Speed and Bearing Data: When speed or bearing data is missing, we estimate the values using the distance and time between consecutive pings. We compute distance using the Haversine distance formula and bearing using the Haversine bearing formula. We discard those approximations where the distance or time between sequential points is infeasible or if the calculated speed exceeds 150km/h.

Identifying Parked Vehicles: We identify parked vehicles

if the speed is below 1 km/h for a time window of 5 minutes. This helps in filtering out data points where vehicles are not actively traveling as to avoid biasing roadways near common parking locations.

Data Cleaning: We remove any duplicated data, data with null locations or timestamps. We also use DBSCAN clustering to detect and remove GPS locations that are abnormal in the context of the rest of the dataset. The large red cluster in fig. 4 is the primary cluster in the Bay Area data collection.



Fig. 4. DBSCAN clustering map of Bay Area RT-GTFS data highlighting a dark red primary cluster surrounded by outlier clusters and points.

C. Merging with OSM

In order to standardize the data and interface with common routing engines, we map incoming data to OSM edges.

1) Finding Matches: We perform a spatial join between OSM edges and GPS coordinates with a tolerance of 20m. This can result in multiple matches based on edge density. To pick the best match a score is determined for each match using the difference between the bearing of the ping and the bearing of the edge as well as the distance from the pings projection onto the edge. The best match is selected using this score and matches under a score threshold are removed.

2) *Removing Sampling Bias:* Sequential GPS pings from the same vehicle on the same edge are aggregated into one data point by averaging speeds to reduce sampling bias during slow periods.

3) Edge-Level Data Aggregation: We aggregate the distribution of speeds for each edge into statistical measures, including standard deviation, mean, and a set of quantiles. The edge level statistical measures are then saved for later use.

V. PROPOSED APPROACH

Our approach, shown in fig. 5, involves encoding each representation of each operational objective into separate instances of the Open Source Routing Machine (OSRM). OSRM is a high-performance routing engine designed for finding the shortest paths in road networks. By using OSRM, we are able to leverage its highly optimized and reliable search methods, including contraction hierarchies and multi-level Dijkstra's algorithm. We then generate a set of initial routes, evaluate, and score them. Once scored we select the recommended route and estimate the travel time.



Fig. 5. Architecture diagram illustrating our approach, where multiple OSRM instances encode different operational objectives for multiple objective routing.

A. Encoding Objectives into OSRM Instances

Each representation of each operational objective, such as road speed, speed variability (represented by multiple quantiles), construction zones, pedestrian encounters, school zones, operator-preferred routes, and legally restricted roadways) is encoded into a separate OSRM instance by updating the internal weights with the edge-level data using osrm-contract.

B. Generating Initial Routes

First, we modify each OSRM instance to generate multiple alternative routes for each query. We then query each objective-specific OSRM instance to generate an initial set of routes. This modification as well as using OSRM's alternative routes parameter, returns several objective optimized paths between the start and end points. The set of generated routes is then consolidated by removing any routes that were duplicated across OSRM instances. We are left with a set of several unique routes between the two points, each optimized relative to differing objectives.

C. Route Evaluation

The consolidated set of initial routes are then re-evaluated using all other objective-specific OSRM instances using highresolution waypoints. The individual *scores* are used in later steps to decide between routes. To ensure metric accuracy, we capture the evaluated routes and compare them to the initial routes. Routes that are not equivalent within a tolerance are indicative of a part of a route deemed impassible in the context

TABLE III Example of User Preference Profile with Time Estimate

Attribute	w	α	β	${\mathcal S}$
Speed	0.6	0.4	0.6	Mean
Construction	0.2	1	0	Mode
Pedestrian	0.2	1	0	Mode
Time Estimate	_	1	2	Mean

of the secondary objective. The resultant metrics from these routes are marked as invalid and dropped.

D. Preference-based Scoring and Optimization

For each proposed route, *scores* generated in the prior step, are used in the route equation, eq. (2), which incorporates a user-configurable preference profile. The preference profile defines the weight w, coefficients α and β , the summary function S for each objective. This allows the operator to balance objectives according to their specific needs. A preference profile with a high weight on speed and speed consistency, but considerations for avoiding construction and pedestrian-prone roads is shown in table III. Table III shows a user preference profile with time estimates for various attributes, users can modify the parameters to optimize the evaluation function for their specific needs, furthermore, our approach supports addition of attributes, for example, preference for scenic routes can be addded to our system.

E. Selecting Route and Estimating Travel Time

The route with the lowest *score*, given by eq. (3), is chosen as the suggested route. The estimated time is calculated using the time estimate field in the user preference profile. An example of a very conservative time estimate is shown in table III.

VI. RESULTS AND DISCUSSIONS

In this section, we evaluate AVATAR and compare the results against several baseline methods. The experiments were conducted on a 32-core Intel Core i9-14900k with an NVIDIA GeForce RTX 4090 (24GB) and 128 GB RAM.

A. Baselines

We evaluate the effectiveness of our proposed method by comparing it against the following baseline routing methods:

Default OSRM: This method uses the standard Open Source Routing Machine (OSRM) without any modifications for AVspecific constraints. This is an off-the-shelf implementation of OSRM v5.28.0, with OpenStreetMap data from Geofabrik.

INRIX Travel Averages: This method uses an OSRM instance with travel speeds updated to be the average INRIX travel speed for a roadway. Only segments which have a corresponding INRIX speed value is updated, otherwise it will use the OSM information.

B. Experimental Setup

We conducted experiments in simulation for single-capacity autonomous vehicles within Nashville, TN, using INRIX data from June 1, 2024 to June 22, 2024. INRIX data was loaded into OSRM instances for both testing and validation to remain consistent. The first 14 days were used in the setting up the AV routing framework as well as the INRIX Travel Averages baseline, while the remaining 8 days were used for validation.

AVATAR Setup: The AV routing framework was created using five OSRM instances, for the 5, 25, 50, 75 and 95 percentiles of speeds. To maximize reliability in this experiment α was set to 0, and β was set to 1 (in eq. (2), making our evaluation function:

$$E(r) = \sigma(c(r))$$

To emphasize the configurable nature of AVATAR, some results are also shown for $\alpha = 1, \beta = 5$, and S = mean. This evaluation function becomes:

$$E(r) = \mu(c(r) + 5 \cdot \sigma(c(r)))$$

For our time estimate, we used $\alpha = 1, \beta = 0$ and S = mean. This makes the time estimate function:

$$E(r) = \mu(c(r))$$

Validation: Data from every hour in the eight day validation data was used to create a separate validation instance, totaling 192 validation instances. The goal of each of these instances is to represent a specific snapshot time where we have complete knowledge of road conditions. The travel times for the routes generated by each method will be computed using these instances.



Fig. 6. Percentage error in estimated travel times. Negative values are underestimations in travel times. The spread represents the precision of the estimate with the median representing the accuracy.

1) Route Calculation: We calculated routes between 382 origin-destination pairs throughout Nashville, TN using three different methods: the AV framework, OSRM with INRIX travel averages and default OSRM. These 1,146 proposed routes are each evaluated using the 192 validation engines to obtain the actual travel times.

2) *Evaluation Metrics:* We evaluated the performance of each routing method using the following metrics:

Travel Time: The actual travel time for each route, as obtained from the validation engine. This is measured in Average Travel Time (seconds).

Travel Time Estimation Accuracy: The error of the traveltime estimation given by the routing engines when compared to the actual travel time for each route, as obtained from the validation engine. This is measured in MSE (seconds) and Average Percent Error (%).

Route Reliability: The consistency of travel times across different runs, measured by the standard deviation of travel times across validation road states. This is measured in standard deviation (seconds).

C. Results

The results of our experiments are summarized in table IV. Our proposed AVATAR routing framework outperformed both default OSRM and our averages baseline in terms of travel time estimate mean squared error (MSE) and travel time estimate average percent error. Specifically, the AVATAR framework demonstrated the lowest travel time estimate MSE of 253.8, which is a significant improvement over the averages baseline MSE of 669.7 and the default OSRM MSE of 13063.3.

Additionally, fig. 6 illustrates the distribution of percentage errors in travel time estimates, with AVATAR showing a median near 0% and a notably smaller inter quartile range (IQR) width, highlighting its precision in travel time predictions. The IQR for AVATAR is only 5.1%, compared to 10.1% for the averages baseline and 20.5% for the default OSRM, demonstrating AVATAR's higher accuracy and precision in estimating travel times.

Moreover, AVATAR also achieved the lowest average percent error of -0.13% in comparison to -0.92% for averages and -26.2% for the default. This reaffirms the enhanced accuracy of our proposed framework.

In terms of travel time consistency, AVATAR obtained the lowest route standard deviation at 12.6 seconds, while both the

 TABLE IV

 Summary of Experimental Results (best results are boxed)





Fig. 7. Standard deviation of route travel times for AVATAR and default planners, with AVATAR showing the lowest variability compared to the baselines, indicating more consistent travel times.

averages baseline and default OSRM showed higher variability in travel times with standard deviations of 20.6 and 20.5 seconds, respectively. Their distribution of standard deviations is shown in fig. 7.

Furthermore, although AVATAR experienced a slightly higher average travel time of 384.8 seconds versus 369.8 seconds for the default and 370.5 seconds for the averages, this is an expected result of favoring more reliable routes and a by-product of using a preference profile that gave no weight to mean speed. Relaxing the reliability preference in AVATAR- β = 5 yields a slightly better average travel time while also outperforming baselines in MSE, percent error and route standard deviation. This result emphasizes the more robust consideration of the road network being used to avoid delay prone routes.

Our experiment clearly evidences the superior performance of the AVATAR framework in providing more accurate and reliable route planning by substantially outperforming the default OSRM and averages baseline across key performance indicators.

D. Sensitivity Analysis

In this section, we analyze the sensitivity of our autonomyaware routing framework to changes in the coefficients α , β for different criteria. This analysis helps to understand how varying these parameters affects the returned routes and their performance. Using the same experimental setup described previously, we conducted experiments by varying the β parameter in eq. (2). Results are summarized in fig. 8. As the weight of the standard deviation increases, the reliability of the route improves and the accuracy of travel time estimations is enhanced. Specifically, the mean squared error (MSE) of travel time predictions decreases from approximately 358.5 seconds when standard deviation is not considered at to 253.8 seconds when the weight reaches its maximum value. In parallel, the added time percentage, which represents the



Fig. 8. Sensitivity of AVATAR route choice performance to standard deviation weight. The left y-axis shows the mean squared error (MSE) of estimated travel times (in seconds), while the secondary y-axis displays both the added time percentage and the route standard deviation (in seconds). The x-axis represents the weight of the standard deviation applied, with the value max denoting a relative α value of 0.



Fig. 9. Wait time distributions for AVATAR and default planners, highlighting trade-offs between acceptance rates and delays. AVATAR achieves zero median and first quartile wait times, while the default planner results in higher variability and longer waits.

percent difference from the general travel time to the lowest recorded value, increases gradually by roughly 8.3%. This trend highlights a clear trade-off: while higher weights help in reducing variability and enhancing prediction accuracy, they concurrently introduce a moderate increase in the travel time.

E. Operational Impact

To evaluate the operational impact of AVATAR on transit operations, a simulation was conducted in Yokohama, Japan using 1,000 ride requests over a 12 hour period, served by 20 single capacity vehicles. AVATAR was configured with RT-GTFS data to create 25, 50, and 75 percentile instances, that were used with preference profile $\alpha = 1$, $\beta = 2$ and time

estimate $\alpha = 1$, $\beta = 2$. Ride requests were then accepted or denied and given pickup times using a constraint driven offline VRP solver. This solver was configured using both default OSRM and AVATAR travel time assumptions. These plans were then evaluated using Google Maps Real-Time Directions API to generate actual pickup and drop off times. Delays were allowed to propagate to the next pickup to capture the possible impact of reliable information in planning.

In this simulation, default OSRM accepted 98.5% of requested rides, resulting in a median pickup time of 4.4 minutes and a maximum wait of 58.5 minutes. Meanwhile, AVATAR accepted only 74.2% of ride requests, resulting in a median pickup time of 0 minutes and a maximum wait time of 15.8 minutes. Planning using the oracle information provided by Google Maps resulted in an acceptance rate of 80.9% of requests and no wait times, shown in fig. 9.

VII. CONCLUSION

In this paper, we introduced AVATAR, an autonomy-aware routing framework designed to address the unique challenges faced by autonomous vehicles (AVs) in on-demand micro transit systems. Our approach prioritizes dependable, lowvariance roadways by incorporating multiple operational objectives such as road speed, speed variability, pedestrian encounters, zoning areas, and operator-preferred routes. By leveraging a multi-criteria decision-making process and runtimeconfigurable preference profiles, our framework allows operators to dynamically balance reliability, speed, and a host of other factors according to their needs.

We demonstrated the effectiveness of our framework through data collection and processing of RT-GTFS, AV data, and traffic data. We generated a comprehensive dataset that accurately reflects real-world roadway conditions. Our experimental results, validated using real-world data from Nashville TN, Bay Area, CA and Yokohama Japan, showed significant improvements in routing reliability and performance compared to traditional routing methods. The proposed framework not only enhances early-stage AV deployment but also facilitates future integration with conventional transit vehicles, paving the way for seamless multimodal operations. AVATAR's flexibility ensures that the routing framework can adapt to dynamic conditions and operational requirements.

In conclusion, our autonomy-aware routing framework advances the state of the art in AV routing by focusing on dependable roadway selection, real-time data integration, and robust decision-making. By addressing the critical challenges of AV deployment, our work contributes to the development of safer, more efficient, and sustainable transportation systems.

ACKNOWLEDGMENT

This material is based upon work sponsored by the Nissan Motor Co., Ltd. Any opinions, findings, conclusions, or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of Nissan. The authors would also like to acknowledge the contributions of the OSM maintainers for their invaluable support in maintaining OpenStreetMap data, the Public Transportation Open Data Center (ODPT) for providing access to valuable transit data, and 511 SF Bay's Open Data Portal for their freely accessible traffic and transit information. Results presented in this paper were obtained using the Chameleon Testbed supported by the NSF.

REFERENCES

- D. J. Fagnant and K. Kockelman, "Preparing a nation for autonomous vehicles: opportunities, barriers and policy recommendations," vol. 77, pp. 167–181.
- [2] R. Etminani-Ghasrodashti, R. Ketankumar Patel, S. Kermanshachi, J. Michael Rosenberger, D. Weinreich, and A. Foss, "Integration of shared autonomous vehicles (SAVs) into existing transportation services: A focus group study," vol. 12, p. 100481.
- [3] J. Hwang, W. Li, L. Stough, C. Lee, and K. Turnbull, "A focus group study on the potential of autonomous vehicles as a viable transportation option: Perspectives from people with disabilities and public transit agencies," vol. 70, pp. 260–274.
- [4] M. M. Rana and K. Hossain, "Connected and autonomous vehicles and infrastructures: A literature review," vol. 16, no. 2, pp. 264–284.
- [5] S. Rezwana and N. Lownes, "Interactions and behaviors of pedestrians with autonomous vehicles: A synthesis," vol. 4, no. 3, pp. 722–745, number: 3 Publisher: Multidisciplinary Digital Publishing Institute.
- [6] M.-Y. Zhang, S.-C. Yang, X.-J. Feng, Y.-Y. Chen, J.-Y. Lu, and Y.-G. Cao, "Route planning for autonomous driving based on traffic information via multi-objective optimization," vol. 12, no. 22, p. 11817, number: 22 Publisher: Multidisciplinary Digital Publishing Institute.
- [7] S. A. Bagloee, M. Tavana, M. Asadi, and T. Oliver, "Autonomous vehicles: challenges, opportunities, and future implications for transportation policies," vol. 24, no. 4, pp. 284–303.
- [8] L. Kloeker, G. Joeken, and L. Eckstein, "Economic analysis of smart roadside infrastructure sensors for connected and automated mobility."
- [9] Y. Liao, J. Gil, R. H. M. Pereira, S. Yeh, and V. Verendel, "Disparities in travel times between car and transit: Spatiotemporal patterns in cities," vol. 10, no. 1, p. 4056, publisher: Nature Publishing Group.
- [10] W.-C. Hu, H.-T. Wu, H.-H. Cho, and F.-H. Tseng, "Optimal route planning system for logistics vehicles based on artificial intelligence," *Journal of Internet Technology*, vol. 21, no. 3, pp. 757–764, 2020.
- [11] A.-E. Taha and N. AbuAli, "Route planning considerations for autonomous vehicles," *IEEE Communications Magazine*, vol. 56, no. 10, pp. 78–84, 2018.
- [12] Q. T. Minh, D. P. Tan, H. N. Le Hoang, and M. N. Nhat, "Effective traffic routing for urban transportation capacity and safety enhancement," *IATSS Research*, vol. 46, no. 4, pp. 574–585, 2022.