Transportation
https://doi.org/10.1007/s11116-024-10548-8

Multifunctional lightweight autonomous vehicles: an agent-based study

Naroa Coretti Sanchez[1] · **Kent Larson**[1]

Accepted: 6 October 2024

Abstract
In mobility-on-demand services, the number of vehicles needed is often determined by peak demand during rush hours, leading to prolonged vehicle idle times during off-peak periods. This surplus capacity presents an opportunity for vehicles to perform additional tasks, potentially enhancing system efficiency and reducing the overall number of vehicles needed in cities. Leveraging agent-based modeling, we evaluate the effectiveness of vehicles catering to on-demand rides and food deliveries in two real-life scenarios: Cambridge, MA, USA, and San Sebastian, Gipuzkoa, Spain. The results show that multifunctional behavior can lead to reduced fleet sizes, with context-specific exceptions. Additionally, a strategic dispatching algorithm is introduced that demonstrates reductions in wait times and overall distances traveled. This research contributes to the understanding of the performance of multifunctional fleets in diverse urban contexts, informing the development of sustainable and resource-efficient mobility systems.

Keywords Shared autonomous vehicles · Micro-mobility · Agent based modeling · Mobility on-demand

Introduction

As global climate concerns escalate, nations worldwide are enacting ambitious agendas to curb transportation emissions (Pörtner et al. 2022). Unfortunately, despite these efforts, a concerning trend persists, with many countries witnessing a rise in transportation-related emissions (Jaramillo et al. 2022). Notably, while countries in the European Union and the United States of America (USA) have successfully reduced carbon emissions across various sectors, transport emissions have continued on an upward trajectory over the last three decades (Crippa et al. 2022). Given that cities already contribute to 62% of global emissions

✉ Naroa Coretti Sanchez
naroa@mit.edu

Kent Larson
kll@mit.edu

[1] MIT Media Lab - City Science, Cambridge, MA, USA

Published online: 27 November 2024

and are expected to represent an even larger share in the future, urban mobility emerges as a critical cornerstone in transitioning towards a more sustainable future (Jaramillo et al. 2022).

In the pursuit of innovative solutions fostering sustainable urban mobility, diverse visions have emerged. Some advocate for reducing reliance on cars through the promotion of walkability and the provision of sustainable transportation alternatives (United Nations Human Settlements Programme (UN-Habitat): 2022; Allam et al. 2022; Jacobs 1961). In contrast, others see autonomous driving technology as an answer to current urban mobility challenges, offering possibilities such as shared autonomous electric vehicles (SAEVs) (Narayanan et al. 2020). By bridging these perspectives and focusing on shared lightweight autonomous vehicles, the field of shared autonomous micromobility (SAmM) aligns the aspiration for pedestrian-friendly cities with the opportunities presented by autonomous driving (Sanchez et al. 2020; Coretti Sanchez et al. 2022a, b; Kondor et al. 2021). Initial research in the realm of SAEVs has focused primarily on autonomous cars and taxis, demonstrating the potential of autonomy to enhance shared transportation services in terms of efficiency, convenience, and affordability (Narayanan et al. 2020; Li et al. 2021; Fagnant and Kockelman 2014; Chen et al. 2016). SAmM takes this exploration a step further, recognizing that the benefits of autonomous vehicles can extend beyond cars to lighter modes of transportation, including bicycles, tricycles, or scooters (Sanchez et al. 2020; Lin 2021; Kondor et al. 2021). These lightweight alternatives align with the broader vision of creating human-centric, walkable, and sustainable cities by shifting away from car dependency. As such, SAmM harmonizes technological advancements with the objective of fostering urban spaces that are not only efficient but also environmentally friendly and human-centric.

SAmM vehicles are currently in development by both industry and academia, with projects focused on autonomous bicycles, tricycles, and scooters (Sanchez et al. 2020; Lin 2021; Tortoise: Tortoise Shared Micromobility 2023; Tech Crunch 2019; Manoeva et al. 2022). Currently, there is a growing body of research exploring the fleet-level behavior of these systems, investigating aspects such as performance, service level, and environmental impacts (Coretti Sanchez et al. 2022c, a, b; Kondor et al. 2021; Genua Cerviño et al. 2024). One of the key findings of such studies is that the number of SAmM vehicles needed for a certain demand heavily depends on peak demand during the morning and evening rush hours (Coretti Sanchez et al. 2022c, a). Despite the significant improvement in the utilization rates of autonomous vehicles compared to their nonautonomous counterparts, SAmM systems still experience a substantial idle time of approximately 80–90% (Coretti Sanchez et al. 2022c). This surplus capacity suggests an opportunity for these vehicles to perform additional tasks during off-peak hours, such as delivering packages, food, or groceries. Vehicles that serve multiple functions throughout the day could improve system efficiency and reduce the overall number of vehicles required in urban areas. This, in turn, could create more open space on the streets, fostering a more human-centric and vibrant urban environment.

The aim of this study is to analyze the potential implications of a multifunctional shared lightweight autonomous vehicle fleets with a specific focus on the combination of on-demand rides and food deliveries. Food deliveries have grown remarkably rapidly in the last five years, increasing fourfold between 2018 and 2021 (Ahuja et al. 2021). This surge has prompted various companies to develop autonomous bots for food and grocery deliveries.[1]

[1] https://www.nuro.ai/, https://www.starship.xyz/, https://www.kiwibot.com/, https://refraction.ai/.

Moreover, food deliveries tend to have a demand pattern that complements the peaks from commuting, making this a particularly compelling case study.

To assess the proposed multifunctional SAmM system, we have developed an agent-based model (ABM) that simulates two scenarios: one with single-use, dedicated SAmM fleets for on-demand rides and food deliveries and another with an integrated multifunctional SAmM system. This modeling approach enables us to evaluate the performance and potential gains stemming from the proposed multifunctional behavior. Additionally, we analyze the performance of a strategic dispatching algorithm that seeks to efficiently assign SAmM vehicles to specific trips based on a decentralized economics-based algorithm.

Our model is applied to two real-world case studies: one in Cambridge, MA, USA, and the other in San Sebastian, Gipuzkoa, Spain. The selection of these cities was based on the availability of comprehensive data and their distinctly contrasting cultural and geographical attributes. This dual application aims to explore the adaptability and effectiveness of the proposed multifunctional system in diverse settings, each characterized by unique demand patterns and network structures.

The remainder of the paper is organized as follows: First, Sections "Related work" and "Contributions" provide an overview of related work and outline our contributions to the field. Second, Section "Methodology" offers a comprehensive explanation of the methodology employed in this study, following the Overview, Design concepts, and Details (ODD) protocol. The study's results are presented and analyzed in Section "Results", while the broader implications are discussed in Section "Discussion". Section Limitations and future work" outlines the limitations of this work and provides avenues for future work. Finally, Section "Limitations and future work" closes with some concluding remarks.

Related work

The agent-based modeling of shared autonomous vehicles (SAV) has been a research field for over a decade, with the first publication by Fagnant and Kockelman (2014). The first study that focused on shared autonomous vehicles (SAEVs), instead, was conducted by Chen et al. (2016). Since then, the field has advanced very significantly, particularly from 2019 onward (Wang et al. 2019; Sayarshad and Gao 2020; Calabrò et al. 2022), as highlighted in the comprehensive literature review by Li et al. (2021).

Previous works primarily focused on larger vehicles such as cars and taxis. However, as the field of SAmM has gained increasing relevance, numerous agent-based models have been developed to examine the performance of autonomous micromobility solutions. In particular, several studies have examined the fleet-level behavior of SAmM as a *mobility-on-demand* service, focusing on aspects such as service quality and environmental impacts. For instance, Coretti Sanchez et al. (2022a, 2022c) developed a simulation tool to analyze the performance of shared autonomous bicycles in urban settings, revealing significant reductions in fleet size requirements compared to traditional bicycle-sharing systems. Similarly, Kondor et al. (2021) analyzed the fleet size and utilization of autonomous scooters, examining their performance in covering trips currently served with dockless shared bicycles, and found significant reductions in fleet size requirements. Additionally, Coretti Sanchez et al. (2022b) reported that, compared with current systems, autonomous bicycles could potentially reduce the CO_2, eq emissions per passenger kilometer traveled. Others have explored decentralized algorithms for future scalability and robustness of SAmM operations; Alfeo et al. (2019) studied the use of SAmM for collecting and transporting trash from the street trash cans to the deposits by employing a bioinspired algorithm, while Coretti Sanchez et al. (2023) studied the use of a decentralized algorithm for vehicle rebalancing in SAmM.

On the other hand, the development of autonomous last-mile delivery systems such as *Nuro*, *Starship*, *Kiwibot* or *Refraction AI*,[2] has also led to the development of several agent based models (Poeting et al. 2019; Haas and Friedrich 2017). However, the literature in this field remains relatively scarce (Li et al. 2021). Specifically focusing on lightweight autonomous vehicles for *food delivery* De Capitani Da Vimercate (2019) presented the first study evaluating the use of autonomous sidewalk robots to serve food delivery trips. A subsequent study by Samouh et al. (2020) focused on comparing three scenarios involving drones, ground robots, and a combination of both. More recently, Genua Cerviño et al. (2024) analyzed the performance of lightweight autonomous vehicles for food deliveries in terms of fleet sizes and environmental impacts, comparing them to the current scenario in which deliveries are carried out by cars.

Regarding *multifunctional* services, even within the broader field of SAEVs the review by Li et al. (2021) underscores that there is a gap in research on multifunctional vehicle fleets. The two most closely related papers are a study by Schlenther et al. (2020) investigating the possibility of renting private autonomous vehicles for parcel deliveries, while Romano Alho et al. (2021), proposed using shared (nonautonomous) vehicles for parcel deliveries in addition to working as a mobility-on-demand system. While these studies are relevant to our current investigation, they do not focus on shared autonomous vehicles.

Contributions

This paper addresses the gap identified in the literature by presenting the first study on the performance of multifunctional shared autonomous vehicles. First, we developed an agent-based model (ABM) for assessing the performance of a multifunctional SAmM system. This model enables the simulation of two scenarios: one with single-use, dedicated SAmM fleets for on-demand rides and food deliveries, and the other featuring an integrated multifunctional SAmM system. This modeling approach allows us to evaluate the performance and potential benefits associated with the proposed multifunctional behavior.

Second, we analyze the effectiveness of a strategic dispatching algorithm that efficiently assigns SAmM vehicles to specific trips based on a decentralized economics-based approach. The calibration and performance evaluation of this strategy provides insights into enhancing the operational efficiency in multifunctional SAmM systems. By developing a dual application with case studies for Cambridge, MA, USA, and San Sebastian, Gipuzkoa, Spain we test the adaptability and effectiveness of the proposed system across distinct urban contexts, each characterized by unique demand patterns and network structures.

Collectively, through a quantitative analysis of system performance, the proposal of strategic operational strategies, and the testing of adaptability across different urban contexts, this research contributes to enhancing the current understanding of multifunctional services.

[2] https://www.nuro.ai/, https://www.starship.xyz/, https://www.kiwibot.com/, https://refraction.ai/.

Methodology

The description of the ABM developed for this study is based on the overview, design concepts, and details (ODD) protocol by Grimm et al. (2020). First, Section "Overview" provides an overview of the model by describing high-level information such as its purpose, patterns, and agents. Second, Section "Design concepts" describes the main design concepts included in the paper. Finally, we describe the details of the model in two different sections: Section "Initialization" describes the initialization of the model, while Section "Submodels" provides a more in-depth overview of the most relevant submodels.

Overview

The overall *purpose* of our model is to investigate the potential for multifunctionality within fleets of shared lightweight autonomous vehicles or shared autonomous micromobility (SAmM). Specifically, it aims to evaluate whether, and to what extent, having fleets that can serve on-demand rides as well as food orders in a certain study area can help to reduce overall fleet sizes as compared to having two independent, single-use fleets. Moreover, we study the impact of two different dispatching strategies: a straightforward dispatcher that assigns the closest available vehicle to each trip and an economics-inspired strategic dispatching mechanism that tries to find a more efficient resource allocation.

The general *patterns* that this model studies are based on system performance. Key performance metrics include: (1) the variation in the number of vehicles needed to serve a specific demand and how it varies across different functionality scenarios (i.e., single function vs. multifunctional fleets), and (2) The variation in wait times across different dispatching strategies (i.e., with and without strategic a dispatching mechanism).

The model includes the following *agents*: agents representing the demand (i.e., on-demand ride requests and food orders), agents representing the vehicle fleet (SAmM vehicles), agents representing the infrastructure available to the vehicles (i.e., road network and charging stations), agents representing demand hot-spots, and logger agents that save the necessary information during the simulation runtime. The state *variables* characterizing these entities are listed in Appendix A.1, Table 4.

Regarding the *spatial* and *temporal* extents, the model's spatial extent is and boundaries are based on the modeled cities (Figs. 2 and 3). Vehicles are considered to move along a road network, which is represented as a graph. The model runs in 5-second steps, which is set as a trade-off between having a high time resolution and keeping run times manageable. The time period modeled in each simulation run represents one week in the fleet-sizing and performance evaluation simulations (Sections "Fleet size", "Performance metrics"). In simulations focused on studying dispatching strategies, instead, the modeled period is one day, which is the day with the highest wait times (Section "Strategic dispatching").

Regarding the principal *processes*, the simulation starts with an initialization process that creates the agents: the road network graph, the charging stations, vehicles, riders, and food orders. The static state variables of each agent (e.g., id, location, trip start time) are defined during this process. The simulation is then run in steps. The behavior of riders and package deliveries is defined through a finite state machine (FSM), as detailed in Fig. 5. Vehicles' behavior is also defined by FSMs, as described in Fig. 6. During each simulation step, agents will check whether conditions for switching between states are met. When conditions are met, agents will update their status and change their behavior to that defined by the new status. The details of this behavior can be found in Section "Submodels".

This ABM has been implemented in the GAMA Platform (Grignard et al. 2013; Taillandier et al. 2019), an open-source simulation tool for spatially explicit, multi-level agent-based models. GAMA has been successfully used for mobility-related studies, including simulations of SAmM systems (Grignard et al. 2018; Alfeo et al. 2019; Coretti-Sanchez et al. 2023), affirming its suitability for our study.

Design concepts

In this section, we provide an overview of the model interactions and stochasticity, which are the most relevant design concepts from the perspective of understanding the model's behavior. A description of the additional design concepts defined according to the ODD methodology can be found in Appendix A.2.

Interactions

The main interactions in this model can be observed in Fig. 1. Users requesting on-demand rides or food deliveries interact with vehicles when they make a request. Note that, while Fig. 1 represents autonomous bicycles, this model could be extended to other types of SAmM vehicles including, bicycles, tricycles, scooters or other single-person, low-speed autonomous vehicles. Such request is direct and assumed to be able to occur over any range in distance via a smartphone app. In this process, users also compete with other users for vehicles. In the case of the straightforward dispatcher, the interaction between users is not mediated. When the strategic dispatcher is active, instead, the vehicle acts as a mediator in the competition between users (Section "Submodels"). Vehicles also interact with other vehicles during the recharging process. The capacity of each station is limited, and, as a consequence, vehicles might need to wait for a charger to become available. In this case, the

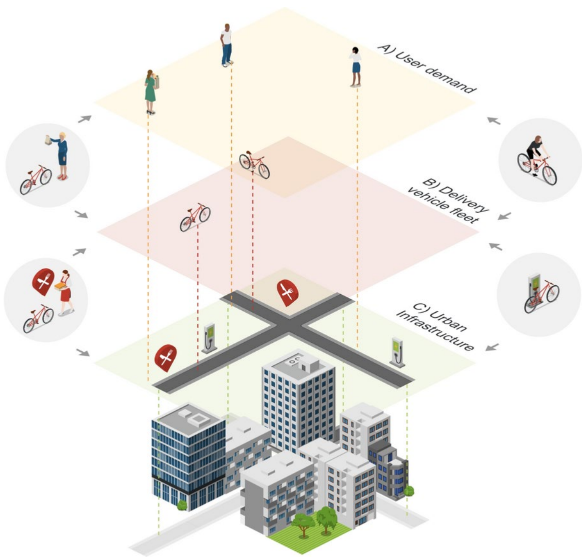

Fig. 1 Multilayer architecture of the model and interactions across layers: **a** urban infrastructure (buildings, roads, charging stations), **b** autonomous micromobility fleet which could include any single-person low-speed autonomous vehicles, and **c** user demand, including rides and food deliveries

interaction is mediated by the charging station, which selects the next vehicle to charge on a first-come first-serve basis.

Stochasticity

Stochasticity is used when initializing the model when selecting where to locate each vehicle and setting the initial battery level. The initial location is set as a random point within the road network, and the battery level is initialized at a random value between the maximum and the minimum. The initial location and battery distribution of vehicles might differ from the natural distribution of vehicles while in use. To mitigate this impact, the simulation is run for one week, and during that process, vehicles travel and relocate naturally. Therefore, the initial location of vehicles and their battery level only have a transitory impact on the simulation model performance. Moreover, each simulation is executed for a repeated number of runs, defined as the number of simulations needed for the output metrics to stabilize and their respective confidence-intervals to fall within the model's requirements (Section "Results").

Initialization

There are different types of data that are imported during the initialization process of the model. This section provides details of the most important initialization datasets, including geospatial, infrastructure, demand, and vehicle-related data.

Geospatial data

The geospatial data for this study includes the boundaries of the selected study area, the buildings within it, and the associated road network. These files are imported in the shapefile format and sourced from OpenStreetMap (OpenStreetMap 2023). As previously discussed, this study focused on two distinct locations: Cambridge, MA, USA, and San Sebastian, Gipuzkoa, Spain (Figs. 2–3), which were selected based on the availability of comprehensive data and their distinctly contrasting network, geographical and cultural attributes. Cambridge has a population of 117,000 inhabitants and an area of $18.40\,\mathrm{km}^2$ while San Sebastian has 186,000 inhabitants and an area of $60.89\,\mathrm{km}^2$. Although San Sebastian has a lower population density, its population is concentrated in specific flat areas near the ocean and rivers. In terms of network structures, Cambridge has a more structured grid-like layout, while San Sebastian's road network is more organic and influenced by topographical elements.

The road network considered for this study is the bicycle network in the case of San Sebastian and the road network for the case of Boston. In the former case, we consider that SAmM vehicles would have lanes that would offer the same coverage as current roads. Before importing the road network into the model, any dangling or duplicate roads are removed and, once imported into the model, the road network is converted into a graph.

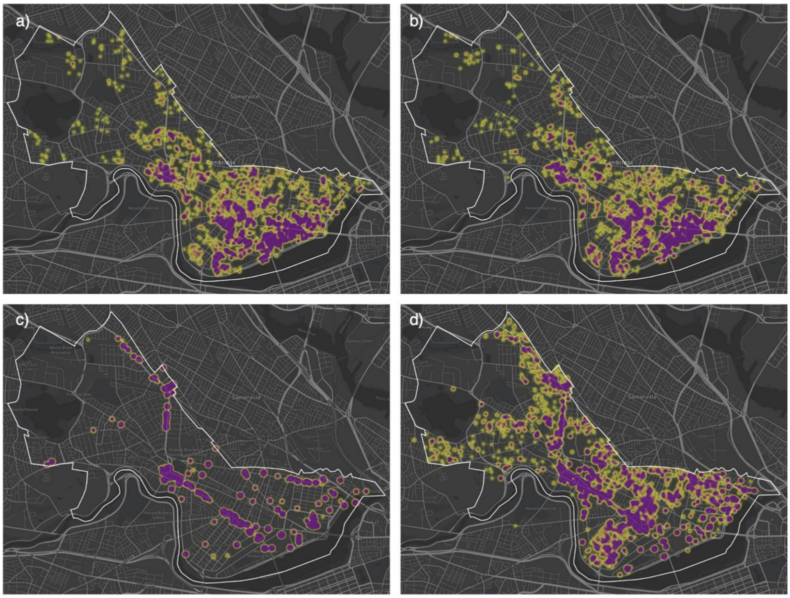

Fig. 2 Heatmaps depicting the demand densities and the boundary of the area considered for the case-study of Cambridge, MA, USA. The areas with the highest density are shown in violet, and the areas with the lower density are shown in yellow. The top row represents on-demand ride **a** origins and **b** destinations. The bottom row represents food order **c** origins and **d** destinations

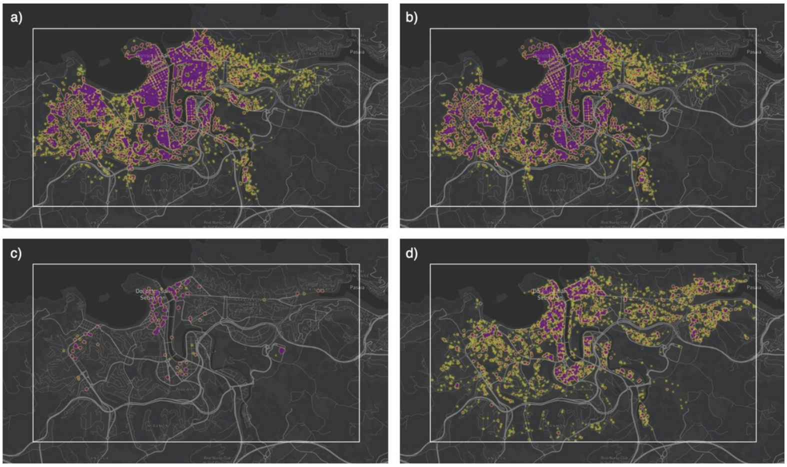

Fig. 3 Heatmaps depicting the demand densities and the boundary of the area considered for the case-study of San Sebastian, Gipuzkoa, Spain. The areas with the highest density are shown in violet, and the areas with lower density are shown in yellow. The top row represents on-demand ride **a** origins and **b** destinations. The bottom row represents food order **c** origins and **d** destinations

Demand for on-demand rides and food deliveries

The model contains two types of demand, on-demand ride and food delivery demand. In the case of Cambridge, on-demand ride demand is based on data from Bluebikes (Bluebikes 2023), the public station-based bicycle-sharing system that operates in Metro Boston. Specifically, we have chosen trips that both begin and end within Cambridge, during the week of October 7–13, 2019, which represents a total of 14,106 trips. In the case of San Sebastian, the on-demand ride demand is based on Dbizi (2023), which is the station-based bicycle sharing system of the City of San Sebastian. We have considered every trip starting and finishing in the study area for a representative week (September 18–24, 2023), which included a total of 19,299 trips. In both cases, the raw data included trip data between stations; therefore, the origins and destinations of the trips were scattered in buildings within a 300 m radius of the stations for a more realistic representation of the demand, following the process described in Coretti Sanchez et al. (2022c).

In terms of food deliveries, we obtained real-world data for San Sebastian from Glovo, an on-demand food delivery operator. This dataset comprises every trip originating and concluding within the study area for the week of September 19–26, 2022, totaling 6499 trips. Conversely, for our study in Cambridge, due to the unavailability of open data on food deliveries, we relied on a high-resolution synthetic database provided for the area (Genua Cerviño et al. 2024). This database, sourced from Replica, Safegraph, and OpenStreetMaps data, consists of two datasets capturing weekday and weekend delivery patterns. To create a representative week of demand, we aggregated five weekdays and two weekend days, resulting in a comprehensive dataset of 30,079 trips over the course of a week.

The last demand-related data used as an input to the model were the hot-spots which represent points around which there is a high concentration of trip start origins. There are two types of demand hot-spots for each case-study: one for user trips and one for food deliveries. The process of generating the dataset has been analogous for both. First, we divided the area into 740 hexagonal cells and calculated the demand density within each cell based on the total count of trip origins within it. Then, we selected the center points of the cells with the highest density. In the case of Cambridge, five hot-spots were chosen for food deliveries and ten for on-demand rides. In the case of San Sebastian, instead, ten hot-spots were chosen for each demand type. These datasets are imported into the model during the initialization process.

Charging stations

In the case of Cambridge, the charging stations are located at the same locations as the current Bluebikes docking stations. The locations of the Bluebikes' docking stations are used as an input to the model. In the case of San Sebastian, the number of stations is also the same as that of the docking stations of Dbizi. However, because food delivery trips have a greater spatial coverage than the stations, the stations were evenly distributed within the study area.

Vehicles

The initial number of vehicles varies across scenarios. In all the scenarios with dynamic fleet sizing, the number of vehicles at the simulation initialization time was defined by the maximum number of concurrent trips at any 15-min period during the week of data (Fig. 4). Specifically, in scenarios in which only food deliveries were considered, the initial number of vehicles was 164 in Cambridge, and 89 in San Sebastian. In scenarios where only user trips were considered, the initial number of vehicles was 86 in Cambridge and 95 in San Sebastian. Finally, the number of vehicles in the multifunctional scenarios was 217 in Cambridge and 122 in San Sebastian. In scenarios without dynamic fleet sizing, the number of vehicles is determined by the results of the scenarios with dynamic fleet sizing.

Submodels

This section contains a comprehensive description of the model's principal submodels. Together, these modules define the behavior of users, individual vehicles, and the vehicle fleet.

User behavior

As shown in Fig. 5, users are initialized in a wandering state. When the current day and time are the same as their departure or food request time, they will request a vehicle. If no vehicles are available within a radius defined by the maximum wait time and the dynamic fleet sizing is active (Section "Submodels"), a new vehicle will be generated at the user's location in the case of on-demand rides and at the restaurant in the case of food deliveries. When the dynamic fleet sizing is not active, the trip will not be served.

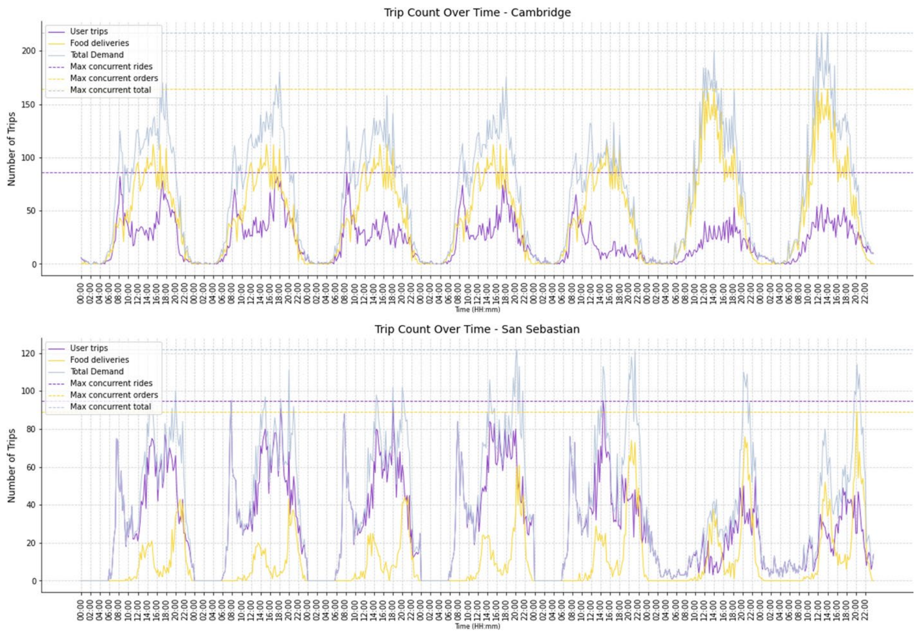

Fig. 4 The evolution of the demand over a one-week period, aggregated by 15-min intervals. The demand data includes on-demand rides, food orders, and the total demand. The horizontal lines represent the maximum number of concurrent trips for each demand type. Top: Cambridge, MA, USA. Bottom: San Sebastian, Gipuzkoa, Spain

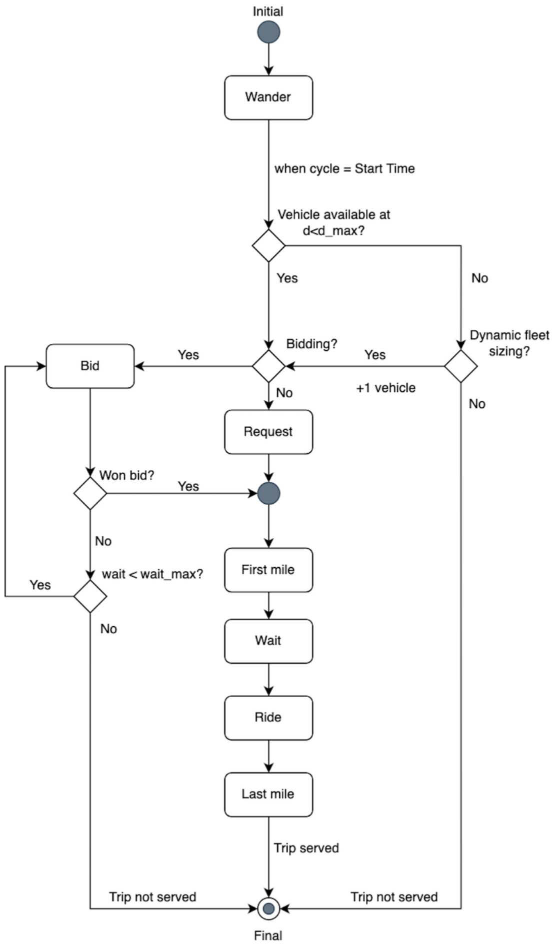

Fig. 5 Diagram depicting the behavior of users requesting on-demand rides or food deliveries as a finite state machine (FSM)

In the scenario with a straightforward dispatcher, the vehicle closest to the demand will be directly assigned to that trip. In the case of the strategic dispatcher being active, the process will follow as described in Section "Submodels". In the case of on-demand rides, once a vehicle is assigned to the trip, the rider will walk to the closest intersection and wait for the vehicle to arrive. Then, the user rides the vehicle to the intersection closest to their destination and walks to their final destination. In the case of food deliveries, the vehicle will drive to pick up the package at the intersection closest to the restaurant. The vehicle will then drive autonomously and transport it to its destination, where the end user will pick it up.

Vehicle behavior

The behavior of the SAmM vehicles is depicted in Fig. 6. Vehicles are initialized in a wandering state. In the scenario with a straightforward dispatcher, whenever a vehicle receives a trip request, the request will be directly assigned to the vehicle. If strategic dispatching is active, the vehicle will follow the process detailed in Section "Submodels". Once a vehicle has a trip assigned, the pick-up location will become its new target and it will drive autono-

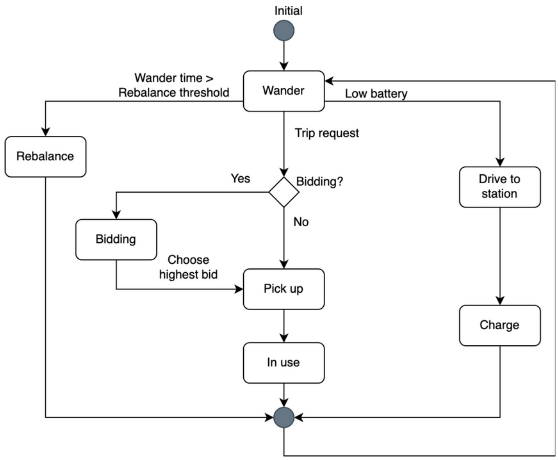

Fig. 6 Diagram depicting vehicle behavior as a finite state machine (FSM)

mously to it. Once there, the rider or food order will be picked up. While in use, the destination location of the trip will become the vehicle's target. In the case of on-demand rides, the user rides the vehicle to the target location while, in the case of food deliveries, the vehicle drives autonomously.

Once the vehicle reaches the destination of the trip, it will become available again, transitioning back into a wandering state. It is important to note that a vehicle that has just completed an on-demand ride may subsequently be assigned to a food delivery, and vice versa. For the purposes of this study, we assume that vehicles are equipped with a separate compartment for food deliveries, eliminating the need for physical transformations between service types.

When a vehicle's battery is below the specified threshold, the charging process will be activated following the process detailed in Section "Submodels". Once charged, vehicles become available again. When a vehicle spends more time wandering that the rebalancing threshold, it transitions to a rebalancing state following the process described in Section "Submodels". Once the rebalancing trip is finished, vehicles transition back into a wandering state.

Moving

At every timestep, vehicles check whether their target is different from their current location; if this is the case, they will drive to it. Vehicles travel along the road network, and they choose the shortest path that connects their trip origin and destination points. After the drive, vehicles will update their remaining battery life by reducing it according to the distance traveled. Since the total battery capacity of the vehicles is defined in kilometers, the battery is reduced by one unit for every kilometer the vehicle travels. If the vehicle is being used by a rider, the vehicle will be in non-autonomous mode and the traveling speed will be 10.2 km/h, which is the average biking speed reported by Jensen et al. (2010). In the remaining cases, the vehicle travels autonomously, and its driving speed is 8 km/h (Coretti Sanchez et al. 2022c). On-demand ride users are assumed to walk, and food deliveries are assumed to be transported by restaurant employees to the closest road intersection at a speed of 5 km/h (Browning et al. 2006).

Charging

Vehicles are assumed to have a battery range of 70 km, based on the e-bike ranges reported by Alli et al. (2010) and the average e-bike battery capacity reported by the International Transport Forum (ITF) (2020). In practice, users of micromobility modes have the flexibility to travel as far as they choose without specifying the trip's endpoint in advance. Consequently, the exact trip length is unknown, and we cannot verify whether the vehicle will have sufficient battery for the entire trip. To manage this uncertainty, we have implemented a protocol where vehicles initiate a recharging sequence if their battery level drops below 25% by the end of any trip. When a vehicle needs to recharge, it will choose its closest charging station as its new target and will drive autonomously to it. Charging stations have a limited number of spaces and vehicles are charged on a first-come first-serve basis. If a station is full, vehicles will queue until a space becomes available. For this study, we assume that the charging stations have 16 chargers, which is the average number of docks at the current Bluebikes stations. Additionally, we consider that the the charging process involves a battery swapping process with a duration of 111 s, which is the average of the two scenarios reported by Huang (2020). Once charged, vehicles become available again. If the rebalancing process is not active, the vehicle will remain idle at the charging station until a trip request is made. However, if rebalancing is active and the corresponding criteria are met, the vehicle will proceed to the nearest demand hot-spot as determined by the rebalancing algorithm.

Dynamic fleet sizing

The goal behind the dynamic fleet-sizing model is to determine the number of vehicles needed to meet the demand in the system at required service level. This service level is defined as a wait time less than 15 min for on-demand rides, and 40 min for food deliveries (Herrmann et al. 2014; US-Foods 2019). Note that these upper limits serve as a conservative threshold to accommodate variability in wait times, which often follow a distribution with a long tail. While the average wait times for on-demand rides are expected to be relatively low compared to this threshold, the 15-min upper limit ensures that the model can account for infrequent yet possible longer waits without oversizing the fleet. Moreover, in the case of food deliveries, this time excludes the time needed to prepare the food. Instead, we followed the criteria in Genua Cerviño et al. (2024) and assumed that the food is ready for when the vehicle arrives to pick it up. Simulations in which dynamic fleet sizing is active start with the number of vehicles determined by the number of concurrent trips (Section "Initialization"). Then, during the simulation runtime, if a trip request cannot be met—either because there are no vehicles available or because the available vehicles are too far relative to the specified threshold—a new vehicle will be generated at the location where the trip was requested (Section "Submodels"). This process is similar to that described by Sheppard et al. (2019) and Zhang and Guhathakurta (2021). New vehicles are generated with the same characteristics as those created during the simulation initialization process.

Rebalancing

Rebalancing is an operation in shared vehicle fleets in which idling vehicles are routed toward high-demand areas to compensate for the gaps between demand and vehicle availability. With this model, we implement a simple rebalancing algorithm in order to maintain tractability. Vehicles that have spent more than the rebalancing threshold time idling—which is set to 12 h—will transition into a rebalancing state. In this state, vehicles choose the closest demand hot-spot as a target, based on the type of demand being studied. In multifunctional scenarios, each vehicle will choose the closest on-demand ride or user hot-spot. In Cambridge, no rebalancing is considered for on-demand rides since, due to the greater symmetry between trip start destinations and trip start origins, the system exhibited a stable performance without rebalancing.

Strategic dispatching

Vehicle dispatching is the task of allocating vehicles to trips. The goal of the strategic dispatching submodel is to efficiently allocate vehicles, which are a scarce resource. For this task, we propose an economics-based algorithm that emulates a bidding process (Wooldridge 2009). One of the main benefits of such algorithms is that they are decentralized and therefore can scale well for large vehicle fleets.

When a vehicle receives a trip request, it will open a bidding process of a predetermined duration. During that time, it will receive bids from other on-demand rides or food deliveries that might want to request its service. The vehicle tracks the value of the highest bid, and the ID of the highest bidder. Once the bidding process is closed, the trip with the highest bidding score is selected. The users will search for another available vehicle and bid again while keeping track of the time spent on this process. If they reach the maximum wait time without winning a bid, the trip will be considered unserved.

The bidding score is calculated using a linear equation that balances demand urgency, wait time, and proximity. Specifically, the bidding score *BS* is calculated using the following equation:

$$BS = w_{urgency} * Urgency + w_{wait} * WaitTime - w_{proximity} * Proximity \tag{1}$$

where $w_{urgency}$, w_{wait}, and $w_{proximity}$ are weighting coefficients to be adjusted based on the relative importance of each factor. Any combination of such weights should always fulfill $w_{urgency} + w_{wait} + w_{proximity} = 1$. The parameter *Urgency* is a binary variable that quantifies the priority level of the demand. It is set to 0 for food delivery and 1 for on-demand rides. This distinction reflects the greater priority assigned to on-demand rides, which requires a lower maximum wait time compared to food deliveries. Consequently, *Urgency* helps prioritize requests based on their urgency and expected service times. *WaitTime* is the time the demand has been waiting for a vehicle assignment, and it has a dynamic value. This parameter is normalized between 0 and 1, where 0 indicates no wait time, and 1 indicates the longest wait time based on the previously defined wait thresholds. This term ensures that customers who have been waiting longer are given a higher priority. Finally, *Proximity* is a measure of how close the demand is to the vehicle. The proximity is calculated as a normalized distance, with values where 0 indicates the closest point and 1 the furthest point, which is based on the radius defined by the maximum wait times and the average traveling speed. Although a customer's proximity to the vehicle remains constant during the bidding process, new requests might emerge that are closer to the vehicle's location.The metric considered to adjust the weights $w_{urgency}$, w_{wait}, and $w_{proximity}$ was the average relative improvement ratio (RIR), which compares the wait times with strategic dispatching to those with the straightforward dispatcher. The goal will be to maximize $Avg_{RIR} = (RIR_r + RIR_f)/2$ where $RIR_r = (W_{SD,r} - W_{nSD,r})/W_{nSD,r}$ and $RIR_f = (W_{SD,f} - W_{nSD,f})/W_{nSD,f}$. The parameters $W_{nSD,r}$ and $W_{nSD,f}$ are the wait times in the baseline scenario with the straightforward dispatching for on-demand rides (*r*) and food deliveries (*f*), respectively, and $W_{SD,r}$ and $W_{SD,f}$ are their relative counterparts in the strategic dispatching scenario.

Results

This section consolidates the outcomes derived from the agent-based model and the experiments described in the preceding sections. The results are systematically analyzed for both the Cambridge and San Sebastian case studies, underscoring the distinct performance of multifunctional services across diverse urban settings.

The section is structured as follows: First, Section "Fleet size" determines the fleet size required to meet demand requirements in each of the studied scenarios, highlighting how it varies between single-function dedicated fleets and multifunctional fleets. Second, Section "Strategic dispatching" illustrates the calibration results of the strategic dispatching algorithm, shedding light on its efficacy in optimizing vehicle assignments for enhanced operational efficiency. Finally, Section "Performance metrics" gathers the main performance metrics in steady-state conditions, incorporating insights from both the fleet-sizing analysis in Section "Fleet size" and the fine-tuned strategic dispatcher outlined in Section "Strategic dispatching"

Fleet size

The fleet size required for each scenario was calculated by following the dynamic fleet-sizing process described in Section "Submodels". To account for stochasticity, each simulation was run 38 times, which was the number of runs needed for the average fleet size between runs to stabilize, and for the 95% confidence interval to be in the range of ± 2 vehicles for all scenarios.

The results of the Cambridge case-study are summarized in Table 1, with supporting graphics in Appendix A.3, Fig. 9 of. The results show that for dedicated fleets, 183 vehicles were needed for on-demand rides and 248 for food deliveries, resulting in a total of **431** vehicles. In contrast, a multifunctional service requires **337** vehicles, resulting in a 21.8% reduction in the number of vehicles needed in the system.

The results of San Sebastian case-study can be found in Table 1 and Appendix A.3, Fig. 12. In this case, 258 vehicles were needed for on-demand rides, and 178 for food deliveries, resulting in a total of **436** vehicles, whereas a multifunctional service needed **476** vehicles; 8.9% more than a single-use service. This increase, rather than a decrease in the number of vehicles needed, highlights the importance of the dual case study conducted.

Table 1 Results of the fleet-sizing process, indicating the number of vehicles required in the Cambridge and San Sebastian case-studies, for the different studied scenarios

Case-study	Cambridge			San Sebastian	
Max. wait time OD-R	15 min	10 min	5 min	15 min	20 min
Single use fleet, OD-R	183	336	754	258	192
Single use fleet, FD	248	248	248	178	178
Single use fleet, Total	431	584	1002	436	370
Multifunctional fleet	337	547	1708	475	311
Fleet-size difference	−21.81%	−6.33%	70.46%	8.94%	−15.95%

The scenarios include varying thresholds for the maximum wait time for on-demand rides (OD-R). Fleet sizes are calculated for single use OD-R and food delivery (FD), as well as for multifunctional services. The bottom row indicates the difference in fleet-size between having two single-use fleets and having a multifunctional service

Further exploration was conducted with the goal of understanding the difference in performance between these case studies, which revealed vehicle deficit concentrations for on-demand rides in the Old Town of San Sebastian, which—being surrounded by ocean, a river, and a hill—is an area with lower network connectivity. To test the potential importance of the network connectivity, we relaxed the maximum wait time for on-demand rides from 15 to 20 min. With a 20 min maximum wait time threshold, the number of vehicles needed in a multifunctional service was 15.95% lower than that needed in single-use fleets (Table 1, Fig. 13 in Appendix A.3).

This led us to investigate whether the Cambridge case-study had a similar inflection point in the maximum wait time that would cause multifunctional services to be less efficient than single-function fleets. With a more restrictive 10-min maximum wait time, a multifunctional service still reduced the number of vehicles needed by 6.34% (Appendix A.3, Fig. 10). However, with a 5-min wait time threshold, the multifunctional service ceased to be efficient, and the fleet size needed increased by 70.46% (Appendix A.3, Fig. 11).

These results indicate that every city and demand-pattern combination has an inflection point in their demand-timeliness requirements that causes a multifunctional service to be less efficient than having dedicated fleets. Such inflection point seems to be case specific and related to the network characteristics of the city. Intuitively, the worst-case scenario for multiple functionalities is one in which the symmetries inherent to single-use modes due to round-trips are broken. In the cases in which the maximum wait time is more flexible, the boundary in which users can look for available vehicles is larger, and, as a consequence, the system becomes less sensitive to small variations in vehicle placement.

According to these results, less time-sensitive demands may be more suitable for multifunctional services. Alternatively, interventions, such as discounts or incentives for users willing to wait longer, could be explored to enhance the overall effectiveness of multifunctional fleets. While these findings contribute to a growing understanding of the contextual factors influencing the success of multifunctionality, they also highlight the need for further research to characterize the specific conditions that determine the viability in the performance of multifunctionality in different scenarios.

In the following subsections, the fleet sizes considered in the case-study of Cambridge will be those of the 15-min threshold scenario, whereas for San Sebastian, we will consider the fleet sizes associated with the 20-min threshold scenario, as determined by their performance inflection points.

Strategic dispatching

This section delves into the performance evaluation of the strategic dispatching algorithm outlined in Section "Submodels" across various parameter configurations. The algorithm's parameters include *MaxBiddingTime*, which represents the duration of the bidding process and the weights $w_{urgency}$, w_{wait}, and $w_{proximity}$ which are the coefficients adjusting the significance of the request type, accumulated user wait time, and proximity of the demand to the vehicle receiving the request, respectively. To calibrate these weights, we employed an exhaustive exploration approach, evaluating all possible parameter combinations. With three parameters, each with five possible values, we assessed a total of fifteen scenarios. For each parameter combination, 15 identical simulation runs were executed to account for stochasticity, and the average wait time was computed across these runs.

Parameter calibration was focused on the day with the highest wait times for the fleet size specified in Section "Fleet size". The results for Cambridge can be found in Fig. 7, with an extended tabular version in Appendix A.4, Table 5. Equivalent results for San Sebastian are available in Figure 8, with an extended tabular version in Appendix A.4, Table 6.

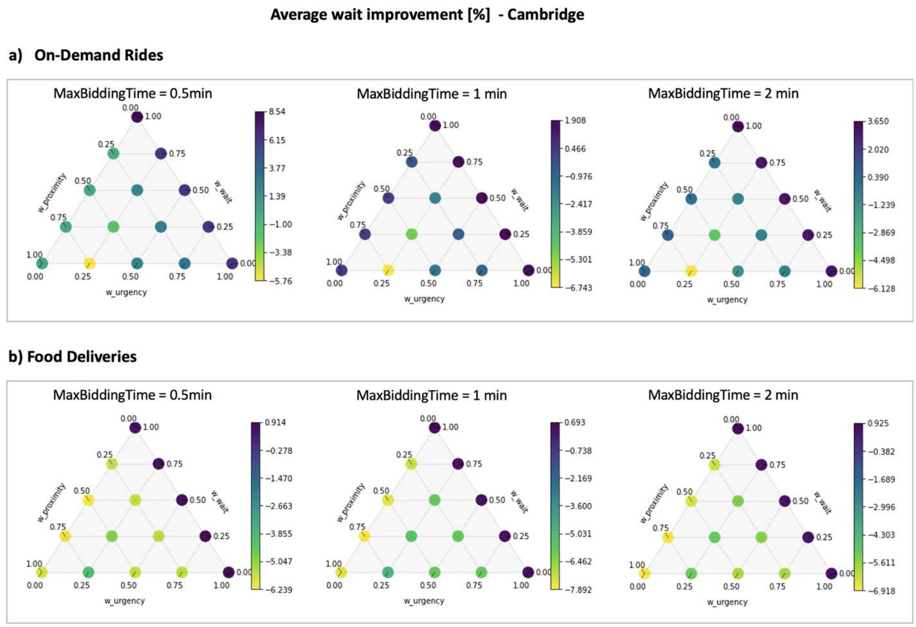

Fig. 7 Percentage improvement in average wait time attributed to the strategic dispatching algorithm in the multifunctional service scenario, compared to the scenario without strategic dispatching. Case-study of Cambridge: **a** on-demand rides and **b** food deliveries. Lighter colors indicate greater improvements, while darker colors denote scenarios with no improvement or negative impact. The three axes on each plot correspond to the bidding algorithm's calibration weights $w_{urgency}$, w_{wait}, and $w_{proximity}$, as defined in Section "Submodels". The parameter *MaxBiddingTime* represents the duration of the bidding process

By evaluating the average relative rate of improvement (avg_{RIR}) as defined in Section "Submodels", the optimal parameter set for Cambridge was a *MaxBiddingTime* of $0.5min$, $w_{urgency} = 0.25$, $w_{wait} = 0$, and $w_{proximity} = 0.75$, yielding an Avg_{RIR} of -6.32%. This result suggests that a shorter bidding process, high importance given to proximity, and consideration of request type contribute to improved performance. Conversely, the consideration of the accumulated wait time has a negligible impact.

In the case of San Sebastian, instead, the maximum Avg_{RIR} is -2.68% which corresponds to a *MaxBiddingTime* of $0.5min$, $w_{urgency} = 0$, $w_{wait} = 0.75$, and $w_{proximity} = 0.25$. In this case, the wait time reductions are significantly greater for food deliveries than for on-demand rides, and, overall, the magnitude of the reduction is smaller than that in the Cambridge case study. While both case-studies benefit from a shorter bidding process, the calibration of the relative importance of the weights seems to be scenario dependent. In the case of San Sebastian, the accumulated wait time becomes the predominant factor, followed by the proximity, whereas considering the type of demand being served does not improve the performance.

It is important to note that, in the case presented in this paper, the best parameters were chosen based on the Avg_{RIR} for the day with the highest wait times. However, by following the same process, the weights could also be calibrated to optimize other performance metrics. For instance, the weights could be optimized for the minimization of average wait times of on-demand rides, or even for minimizing the wait time of the users who wait most. The choice of such metrics will ultimately depend on the specific priorities in each context.

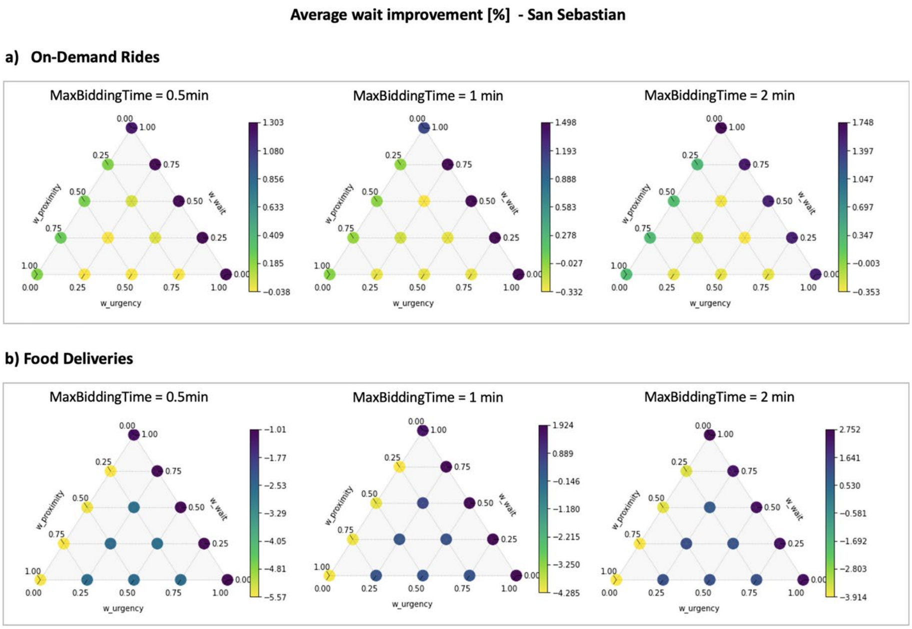

Fig. 8 Percentage improvement in average wait time attributed to the strategic dispatching algorithm in the multifunctional service scenario, compared to the scenario without strategic dispatching. Case-study of San Sebastian: **a** on-demand rides and **b** food deliveries. Lighter colors indicate greater improvements, while darker colors denote scenarios with no improvement or negative impact. The three axes on each plot correspond to the bidding algorithm's calibration weights $w_{urgency}$, w_{wait}, and $w_{proximity}$, as defined in Section "Submodels". The parameter *MaxBiddingTime* represents the duration of the bidding process

Performance metrics

Considering the fleet sizes and calibrated strategic dispatching algorithm described in Sections "Fleet size" and "Strategic dispatching", this section analyzes the performance metrics of the proposed system under these conditions. To obtain the results in this section, and the model was run for the simulated week for 38 times to account for stochasticity, and the numbers shown are the average of these runs.

Table 2 shows the results of the Cambridge case-study. In this case, the number of vehicles needed in the multifunctional scenario is 21.8% lower than that in the single function scenario. The results show that with these fleet sizes, most wait time related metrics improve, except for the average with time for on-demand rides. In terms of distances, there was a 1.1% reduction in the total distance traveled by the system. When introducing the strategic dispatching algorithm, the wait times improve both for on-demand rides and for food deliveries, with more significant reductions for food deliveries; for instance, improvements in average wait time are of 1.0% for on-demand rides, and 9.95% for food deliveries. In addition to reducing the wait times, the strategic dispatching also affects the total distance traveled by the system, reducing it by an additional 0.96% compared to the multifunctional scenario with no strategic dispatching, and by 2.06% compared to the single-function scenario. Additional metrics on the time and distance distributions per task can be found in Appendix A.5, Fig. 14.

The equivalent results for the San Sebastian case-study can be found in Table 3. For the 20-minute maximum wait time threshold, the number of vehicles needed in the multifunctional scenario is 15.95% lower than that in the scenario with single-function dedicated fleets. With these fleet sizes, in the multifunctional fleet scenario with no strategic dispatching, the wait times are reduced for food deliveries but increase for on-demand rides. The results also showed that the total distance traveled by the system decreased by 5.45%. Wait times decrease significantly when the strategic dispatching is implemented, with reductions of 2.7% and 18.54% in the average wait times for on-demand rides and food-deliveries respectively. Strategic dispatching also reduces the total distance traveled by the system by 3.23% compared to the multifunctional scenario with no strategic dispatching, and by

Table 2 Absolute and relative performance metrics for the different scenarios in the Cambridge, USA, case-study: single use fleets, multifunctional service, and multifunctional service with the strategic-dispatching algorithm (S.D.)

	Single use fleets		Multifunctional service		Multif. service with S.D		Difference w/ and w/o S.D.	
	OD-R	FD	OD-R	FD	OD-R	FD	OD-R	FD
$N^{\underline{o}}$ vehicles [-]	431		337		337		–	
Avg_wait [min]	3.68	4.85	3.80	4.40	3.76	3.97	–1.01%	–9.95%
Q10_wait [min]	0.92	1.25	1.00	1.25	0.92	1.00	–8.33%	–20.00%
Q90_wait [min]	7.75	9.17	7.33	8.58	7.33	7.83	0.00%	–8.74%
Distance [km]	3447562		3409557		3376510		–	
Dist. reduction	–		–1.10%		–2.06%		–	

These scenarios include on-demand rides (OD-R) and food deliveries (FD). The performance metrics include the average, 10th and 90th deciles for wait times the improvement in these metrics with and without SD in the multifunctional scenario; the total system distances for the simulated period; and their respective changes from the single-use to the multifunctional service scenario

Table 3 Absolute and relative performance metrics for the different scenarios in the San Sebastian, Spain, case-study: single use fleets, multifunctional service, and multifunctional service with the strategic-dispatching algorithm (S.D.)

	Single use fleets		Multifunctional service		Multif. service with S.D		Difference w/ and w/o S.D.	
	OD-R	FD	OD-R	FD	OD-R	FD	OD-R	FD
$N^{\underline{\text{o}}}$ vehicles [-]	370		311		311		–	
Avg_wait [min]	3.83	9.80	4.43	7.71	4.31	6.66	–2.70%	–13.54%
Q10_wait [min]	0.92	1.67	0.92	1.42	0.75	1.17	–18.18%	–17.65%
Q90_wait [min]	8.25	19.25	9.75	17.42	9.67	16.17	–0.85%	–7.18%
Distance [km]	2890286		2732742		2639431		–	
Dist. reduction	–		–5.45%		–8.68%		–	

These scenarios include on-demand rides (OD-R) and food deliveries (FD). The performance metrics include the average, 10th and 90th deciles for wait times the improvement in these metrics with and without SD in the multifunctional scenario; the total system distances for the simulated period; and their respective changes from the single-use to the multifunctional service scenario

8.68% compared to the single-function scenario. Additional metrics on the temporal and distance distributions per task can be found in Appendix A.5, Figure 20.

Discussion

This study uncovers the complex dynamics in multifunctional vehicle fleets, revealing that multifunctional services can deliver significant efficiency gains, particularly in cities with strong network connectivity and flexible demand patterns. However, our findings also identify a context-specific inflection point where the benefits of multifunctionality diminish, indicating that in scenarios with tight wait time constraints, dedicated fleets may outperform multifunctional models.

Moreover, while the strategic dispatching algorithm notably improved wait times and reduced total travel distances, the necessity for context-specific calibration highlights the importance of tailoring the bidding mechanism to the unique demands and characteristics of each scenario, suggesting that a one-size-fits-all approach is unlikely to be optimal.

Overall, this study underscores the potential of multifunctional fleets to achieve substantial efficiency gains, particularly in well-connected urban areas with flexible demand requirements. However, it also emphasizes the critical need for careful consideration of local network characteristics and demand patterns. Future research should further explore the broader implications of these findings, particularly across diverse urban environments, to refine strategies for the implementation and optimization of multifunctional vehicle services.

Limitations and future work

This study has several limitations that should be acknowledged. First of all, the demand considered for this model is that of current bike-sharing systems. This assumption arises from the absence of real-world data specific to SAmM systems. Future research could refine these estimates by exploring how demand may shift with the introduction of autonomy, potentially through dedicated surveys or pilot studies.

Another limitation is the implicit treatment of traffic congestion. We modeled congestion indirectly through vehicle speed definitions, which does not fully capture the variations in traffic conditions that the system might induce. Although our model uses relatively small fleet sizes compared to current vehicle volumes, understanding how congestion impacts performance would be crucial for higher demand scenarios.

Moreover, while our rebalancing algorithm is effective, it is a simplified model that does not account for the number of vehicles available at each hot-spot. Integrating real-time information into the rebalancing algorithm would significantly improve system responsiveness and robustness, particularly in addressing unbalanced demand scenarios.

Lastly, in calibrating the strategic dispatching algorithm we conducted an exhaustive search for the best parameter combinations, but we only evaluated discrete values of each parameter. As a consequence, this approach may have overlooked potential global maxima that could exist between the discrete values studied. Consequently, there may be unexplored parameter settings that could yield even better performance.

Conclusions

The investigation of lightweight, shared autonomous vehicles in this study has provided valuable insights into the potential impacts and limitations of multifunctional services. The analysis of real-world scenarios in Cambridge, MA, USA, and San Sebastian, Gipuzkoa, Spain, contributes significant considerations for advancing sustainable and resource-efficient urban mobility.

The findings underscore a promising avenue for improving efficiency and reducing the overall number of vehicles through the adoption of multifunctionality with strategic dispatching algorithms. Notably, a substantial reduction in fleet sizes, coupled with decreased total distance traveled, suggests tangible gains in resource utilization. However, the results reveal a context-specific dependency, with instances where multifunctional services may require an increase in the number of vehicles. This variability emphasizes the critical role of factors such as road network characteristics, spatio-temporal demand interactions, and the time-sensitivity of demand. These outcomes reflect that the efficiency of multifunctional services is not universally guaranteed, highlighting the necessity for tailored case-studies that capture the intricacies of specific urban contexts.

In summary, this study highlights the potential gains and context-specific nature of the impact of multiple functionalities on fleet performance and the role of strategic dispatching algorithms in enhancing operational efficiency. The context-dependency of the results underscores the need to fine-tune strategies to address the unique characteristics of each urban environment.

While these study focuses on lightweight autonomous vehicles, these obtained results and conclusions might have implications for current on-demand services.

Overall, this research contributes to the understanding of multifunctional fleets and paves the way for their effective development and implementation in diverse urban landscapes.

Appendix

State variables

See Table 4.

Table 4 State variables for all agent types within the simulation model, including variable name, type, and description

Agent	Variable name	Variable type	Meaning
Riders	id	int, static	Unique identification number
	start_time	[day (int), hour (int), minute (int)], static	Time at which the user will start their trip
	start_location	[latitude (float), longitude (float)], static	Coordinates of the trip origin location
	end_location	[latitude (float), longitude (float)], static	Coordinates of the trip destination location
	current_location	[latitude (float), longitude (float)], dynamic	Coordinates of current location
	state	[“wandering”, “requesting_ride”, “bidding”, “first_mile”, “waiting”, “riding”, “last_mile”], dynamic	Current state in which the person is
	vehicle_assigned	pointer to vehicle agent, dynamic	Vehicle that has been assigned to its trip
	queue_time	minute (int), dynamic	Time that the person has spent waiting for a vehicle assignation
Food orders	id	int, static	Unique identification number
	start_time	[day (int), hour (int), minute (int)], static	Time at which the user ordered the delivery
	start_location	[latitude (float), longitude (float)], static	Coordinates of the restaurant (trip origin location)
	end_location	[latitude (float), longitude (float)], static	Coordinates of the trip destination location
	current_location	[latitude (float), longitude (float)], dynamic	Coordinates of current location
	state	[“wandering”, “requesting_ride”, “bidding”, “first_mile”, “waiting”, “riding”, “last_mile”], dynamic	Current state in which the delivery is (See FSM in figure XX)
	vehicle_assigned	pointer to vehicle agent, dynamic	The vehicle that has been assigned for this delivery
	queue_time	minute (int), dynamic	Time that the order has spent waiting for a vehicle assignation
Vehicles	id	int, static	Unique identification number
	location	[latitude (float), longitude (float)], dynamic	Coordinates of current location
	state	[“wandering”, “picking_up”, “in_use”, “drive_to_station”, “charging”, “rebalancing”], dynamic	Current state in which the vehicle is
	battery_life	Int in [0, 100], dynamic	Remaining battery percentage

Table 4 (continued)

Agent	Variable name	Variable type	Meaning
	delivery_assigned	pointer to food order agent, dynamic	The food order package that has been assigned to this vehicle
	rider_assigned	pointer to rider agent, dynamic	The rider that has been assigned to this vehicle
	highest_bid	Int, dynamic	Highest bid received during the current bidding period
	bid_start	[hour (int), minute (int)], dynamic	Time at which bidding process started
Charging stations	id	int, static	Unique identification number
	location	[latitude (float), longitude (float)], static	Coordinates of location
	vehicles_list	[vehicle pointer 1, ..., vehicle pointer m], dynamic	List of vehicles to charge
Demand hot-spots	id	int, static	Unique identification number
	Location	[latitude (float), longitude (float)], static	Coordinates of location
	Type	bool, static	Defines whether it is a food order or rider hot-spot

Design concepts

The design concepts in this model are defined as follows.

- *Emergence* The model's primary result is the fleet size needed for a certain demand under different scenarios. These results emerge from the characteristics of the demand patterns, network structure, vehicle characteristics, and rebalancing processes. The secondary results, which include performance metrics such as wait times or miles traveled, also emerge from these factors, in addition to the fleet size and the dispatching algorithm.
- *Adaptation* In the dynamic fleet-sizing process, users and food deliveries consider vehicles within a proximity threshold, reflecting their maximum acceptable wait time. Users select the nearest available vehicle to minimize wait times and overall trip duration. Strategic dispatching involves users bidding to their closest available vehicle, and vehicles choosing trips based on bidding values. Adaptive behavior extends to vehicle charging, where vehicles choose the nearest charging station to reduce travel distance and battery consumption. Rebalancing occurs when vehicles idle beyond a time threshold, and they reposition themselves to demand hot-spots, adapting to future demand patterns
- *Objectives* Fleet-sizing experiments aim to determine the minimum fleet size required to meet demand with a specified service level. The strategic dispatching algorithm's objective is to efficiently allocate vehicles, which are a scarce resource in the system.
- *Learning and prediction* The model lacks explicit learning processes. While there is no explicit prediction process, either, when vehicles rebalance towards demand hot-spots they do so with the implicit prediction that that is where future demand is likely to occur.
- *Sensing* Users requesting a ride or food delivery are assumed to know the availability

and location of vehicles which allows them to send a request to the closest available one. Similarly, vehicles are assumed to know the location of charging stations and of demand hot-spots. Lastly, both users, and vehicles traveling in autonomous mode are also assumed to know the shortest path that connects any two given locations.

- *Collectives* Three natural collectives exist in the model: users requesting on-demand rides, users requesting food deliveries, and the vehicles constituting the servicing fleet. While not explicitly represented as collectives, these entities interact within the model.
- *Observation* Individual-level observations are collected during simulations, including user trips, food deliveries, and vehicle actions. These observations encompass trip details, vehicle tasks, distances traveled, and battery information. These datasets are crucial for evaluating system performance, as discussed in Section “Results”.

Fleet sizing

See Figs. 9, 10, 11, 12, 13.

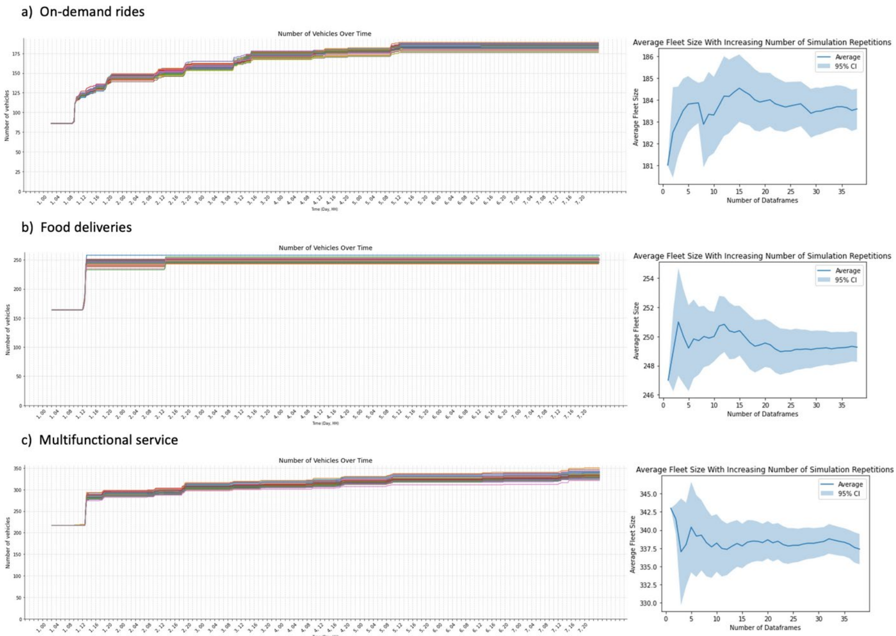

Fig. 9 Cumulative fleet-sizing results following the dynamic fleet sizing process in Section “Submodels” for the case-study of Cambridge, US, with a 15-min maximum wait threshold for on-demand rides. Scenarios: **a** on-demand rides, **b** food deliveries and, c) multifunctional service. Left: Evolution of the dynamic fleet-sizing process in terms of the increase in the number of vehicles during the simulated week. Right: The convergence of the average fleet size and the 95% confidence interval for an increasing number of simulation runs

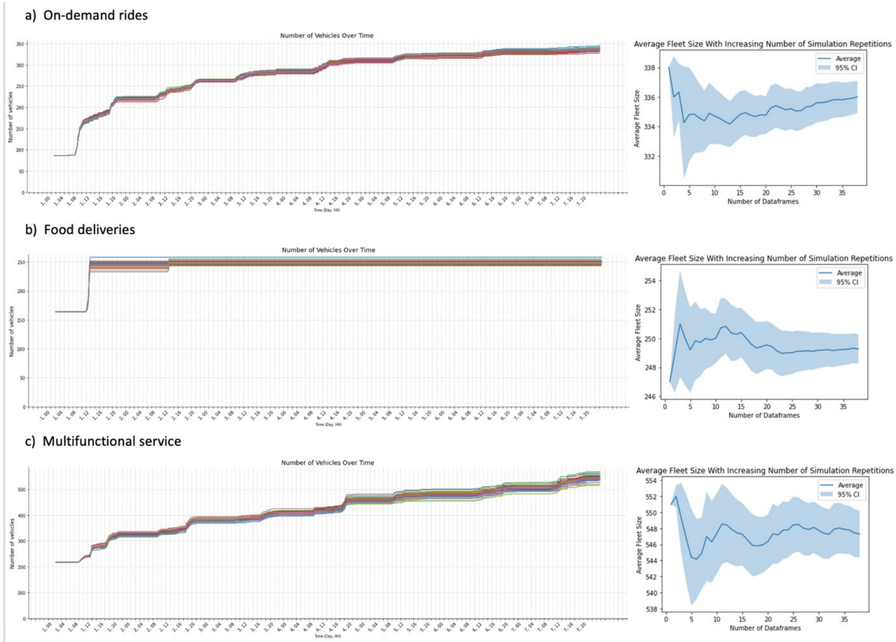

Fig. 10 Cumulative fleet-sizing results following the dynamic fleet sizing process in Section "Submodels" for the case-study of Cambridge, US, with a 10-min maximum wait threshold for on-demand rides. Scenarios: **a** on-demand rides, **b** food deliveries and, **c** multifunctional service. Left: Evolution of the dynamic fleet-sizing process in terms of the increase in the number of vehicles during the simulated week. Right: The convergence of the average fleet size and the 95% confidence interval for an increasing number of simulation runs

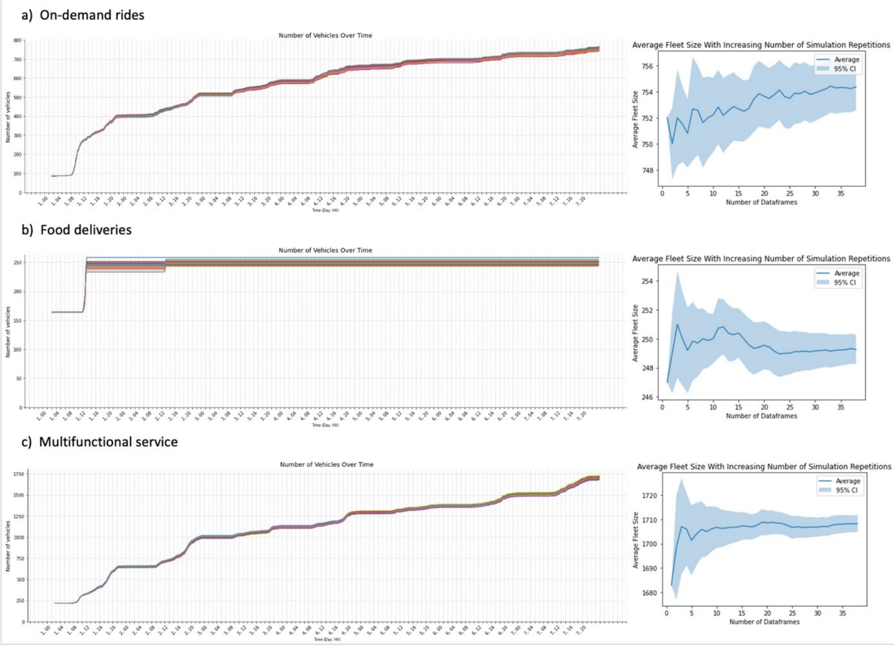

Fig. 11 Cumulative fleet-sizing results following the dynamic fleet sizing process in Section "Submodels" for the case-study of Cambridge, US, with a 5-min maximum wait threshold for on-demand rides. Scenarios: **a** on-demand rides, **b** food deliveries and, **c** multifunctional service. Left: Evolution of the dynamic fleet-sizing process in terms of the increase in the number of vehicles during the simulated week. Right: The convergence of the average fleet size and the 95% confidence interval for an increasing number of simulation runs

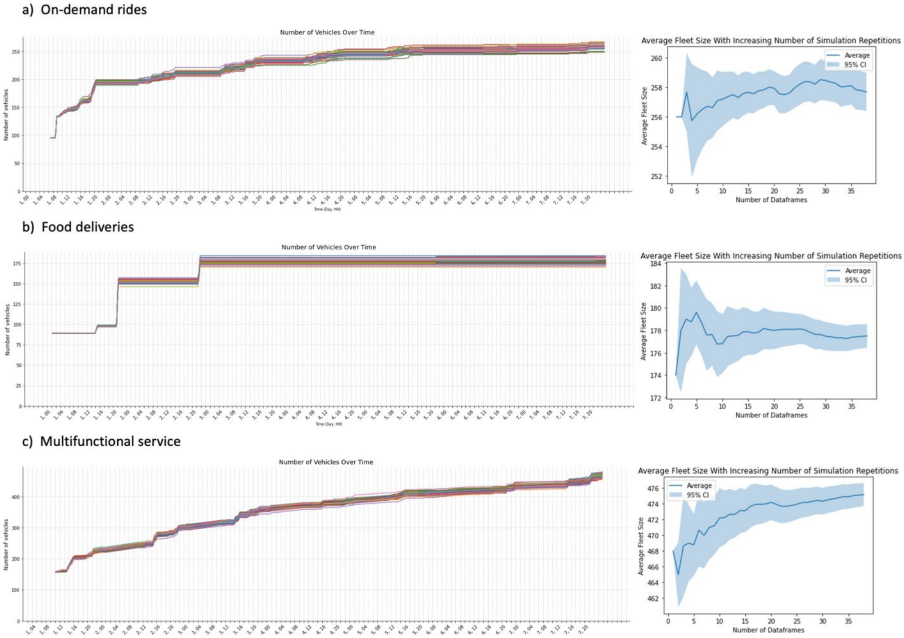

Fig. 12 Cumulative fleet-sizing results following the dynamic fleet sizing process in Section "Submodels" for the case-study of San Sebastian, Spain, with a 15-minute maximum wait threshold for on-demand rides. Scenarios: **a** on-demand rides, **b** food deliveries and, **c** multifunctional service. Left: Evolution of the dynamic fleet-sizing process in terms of the increase in the number of vehicles during the simulated week. Right: The convergence of the average fleet size and the 95% confidence interval for an increasing number of simulation runs

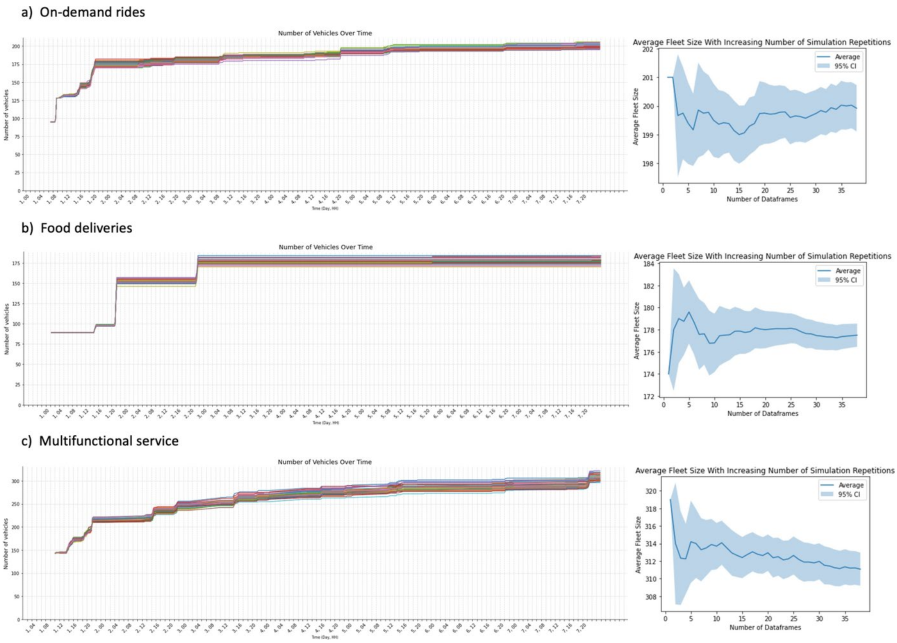

Fig. 13 Cumulative fleet-sizing results following the dynamic fleet sizing process in Section "Submodels" for the case-study of San Sebastian, Spain, with a 20-min maximum wait threshold for on-demand rides. Scenarios: **a** on-demand rides, **b** food deliveries and, **c** multifunctional service. Left: Evolution of the dynamic fleet-sizing process in terms of the increase in the number of vehicles during the simulated week. Right: The convergence of the average fleet size and the 95% confidence interval for an increasing number of simulation runs

Strategic dispatching

See Tables 5 and 6

Table 5 Performance of the multifunctional service with the strategic dispatching in Cambridge, US

Strategic dispatching metrics				On-demand ride metrics		Food delivery metrics		Avg RIR
MaxBiddingTime	w_urgency	w_wait	w_proximity	Avg_ wait [min]	% Impr. Avg	Avg_ wait [min]	% Impr. Avg	Avg RIR _ Avg
0	0	0	0	4.26	–	4.46	–	–
0.5	0	0	1	4.23	−0.60%	4.13	−7.43%	−4.02%
0.5	0	0.25	0.75	4.26	0.07%	4.12	−7.46%	−3.70%
0.5	0	0.5	0.5	4.23	−0.72%	4.13	−7.29%	−4.00%
0.5	0	0.75	0.25	4.19	−1.67%	4.15	−6.93%	−4.30%
0.5	0	1	0	4.24	−0.37%	4.46	0.20%	−0.09%
0.5	**0.25**	**0**	**0.75**	**3.92**	**−7.83%**	**4.24**	**−4.81%**	**−6.32%**
0.5	0.25	0.25	0.5	4.02	−5.48%	4.22	−5.19%	−5.33%
0.5	0.25	0.5	0.25	4.14	−2.69%	4.2	−5.85%	−4.27%
0.5	0.25	0.75	0	4.26	0.00%	4.46	0.00%	0.00%
0.5	0.5	0	0.5	4.14	−2.73%	4.2	−5.84%	−4.29%
0.5	0.5	0.25	0.25	4.17	−2.17%	4.2	−5.78%	−3.97%
0.5	0.5	0.5	0	4.26	0.00%	4.46	0.00%	0.00%
0.5	0.75	0	0.25	4.15	−2.50%	4.19	−5.93%	−4.22%
0.5	0.75	0.25	0	4.26	0.00%	4.46	0.00%	0.00%
0.5	1	0	0	4.26	0.00%	4.46	0.00%	0.00%
1	0	0	1	4.28	0.44%	4.12	−7.59%	−3.57%
1	0	0.25	0.75	4.27	0.32%	4.1	−7.89%	−3.79%
1	0	0.5	0.5	4.27	0.33%	4.12	−7.60%	−3.63%
1	0	0.75	0.25	4.24	−0.39%	4.13	−7.28%	−3.84%
1	0	1	0	4.33	1.74%	4.49	0.69%	1.22%
1	0.25	0	0.75	3.97	−6.74%	4.24	−4.92%	−5.83%
1	0.25	0.25	0.5	4.05	−4.90%	4.2	−5.65%	−5.27%
1	0.25	0.5	0.25	4.18	−1.93%	4.2	−5.85%	−3.89%
1	0.25	0.75	0	4.34	1.91%	4.48	0.46%	1.19%
1	0.5	0	0.5	4.19	−1.51%	4.19	−5.89%	−3.70%
1	0.5	0.25	0.25	4.23	−0.75%	4.2	−5.81%	−3.28%
1	0.5	0.5	0	4.34	1.91%	4.48	0.46%	1.19%
1	0.75	0	0.25	4.22	−0.90%	4.19	−5.87%	−3.38%
1	0.75	0.25	0	4.34	1.91%	4.48	0.46%	1.19%
1	1	0	0	4.34	1.89%	4.48	0.45%	1.17%
2	0	0	1	4.27	0.32%	4.15	−6.83%	−3.25%
2	0	0.25	0.75	4.29	0.70%	4.15	−6.92%	−3.11%
2	0	0.5	0.5	4.26	0.12%	4.17	−6.50%	−3.19%
2	0	0.75	0.25	4.25	−0.18%	4.17	−6.38%	−3.28%
2	0	1	0	4.41	3.65%	4.5	0.92%	2.29%
2	0.25	0	0.75	4	−6.13%	4.23	−5.15%	−5.64%

Table 5 (continued)

Strategic dispatching metrics				On-demand ride metrics		Food delivery metrics		Avg RIR
MaxBiddingTime	w_urgency	w_wait	w_proximity	Avg_ wait [min]	% Impr. Avg	Avg_ wait [min]	% Impr. Avg	Avg RIR _ Avg
2	0.25	0.25	0.5	4.1	–3.76%	4.23	–5.12%	–4.44%
2	0.25	0.5	0.25	4.22	–0.78%	4.22	–5.41%	–3.09%
2	0.25	0.75	0	4.39	3.13%	4.48	0.62%	1.87%
2	0.5	0	0.5	4.2	–1.27%	4.19	–5.97%	–3.62%
2	0.5	0.25	0.25	4.23	–0.65%	4.21	–5.52%	–3.09%
2	0.5	0.5	0	4.39	3.13%	4.48	0.62%	1.87%
2	0.75	0	0.25	4.23	–0.55%	4.19	–5.94%	–3.25%
2	0.75	0.25	0	4.39	3.13%	4.48	0.62%	1.87%
2	1	0	0	4.4	3.23%	4.48	0.63%	1.93%

Each row represents a certain parametrization of the strategic dispatching algorithm which includes the duration of the bidding process, *MaxBiddingTime*, as well as the calibration weights, $w_{urgency}$, w_{wait}, $w_{proximity}$. Columns show the average wait time for on-demand rides and food deliveries within that scenario, and how it compares with the scenario without strategic dispatching. The last column represents the average Relative Improvement Ratio (Avg_{RIR}) as defined in Section "Submodels". The row in bold font represents the highest Avg_{RIR}

Table 6 Performance of the multifunctional service with the strategic dispatching in San Sebastian, Spain

Strategic dispatching metrics				On-demand ride metrics		Food delivery metrics		Avg RIR
MaxBiddingTime	w_urgency	w_wait	w_proximity	Avg_ wait [min]	% Impr. Avg	Avg_ wait [min]	% Impr. Avg	Avg RIR _ Avg
0	0	0	0	3.53	–	6.27	–	–
0.5	0	0	1	3.53	0.21%	5.92	-5.-51%	-2.65%
0.5	0	0.25	0.75	3.53	0.26%	5.92	-5.54%	-2.64%
0.5	0	0.5	0.5	3.53	0.23%	5.92	-5.48%	-2.62%
0.5	**0**	**0.75**	**0.25**	**3.53**	**0.21%**	**5.92**	**-5.57%**	**-2.68%**
0.5	0	1	0	3.57	1.22%	6.19	-1.16%	0.03%
0.5	0.25	0	0.75	3.52	-0.04%	6.1	-2.62%	-1.33%
0.5	0.25	0.25	0.5	3.52	-0.03%	6.1	-2.61%	-1.32%
0.5	0.25	0.5	0.25	3.53	0.04%	6.1	-2.69%	-1.32%
0.5	0.25	0.75	0	3.57	1.30%	6.2	-1.01%	0.15%
0.5	0.5	0	0.5	3.52	-0.04%	6.1	-2.62%	-1.33%
0.5	0.5	0.25	0.25	3.53	0.02%	6.1	-2.66%	-1.32%
0.5	0.5	0.5	0	3.57	1.30%	6.2	-1.01%	0.15%
0.5	0.75	0	0.25	3.52	-0.04%	6.1	-2.62%	-1.33%
0.5	0.75	0.25	0	3.57	1.30%	6.2	-1.01%	0.15%
0.5	1	0	0	3.57	1.30%	6.2	-1.01%	0.15%
1	0	0	1	3.52	-0.01%	6.01	-4.17%	-2.09%
1	0	0.25	0.75	3.52	-0.05%	6.01	-4.15%	-2.10%
1	0	0.5	0.5	3.52	-0.02%	6.01	-4.07%	-2.04%
1	0	0.75	0.25	3.52	-0.03%	6	-4.28%	-2.16%
1	0	1	0	3.56	1.09%	6.37	1.63%	1.36%
1	0.25	0	0.75	3.52	-0.28%	6.29	0.29%	0.01%
1	0.25	0.25	0.5	3.52	-0.22%	6.28	0.23%	0.01%
1	0.25	0.5	0.25	3.51	-0.33%	6.3	0.56%	0.11%
1	0.25	0.75	0	3.58	1.50%	6.39	1.92%	1.71%
1	0.5	0	0.5	3.52	-0.28%	6.29	0.29%	0.01%
1	0.5	0.25	0.25	3.52	-0.27%	6.29	0.34%	0.03%
1	0.5	0.5	0	3.58	1.50%	6.39	1.92%	1.71%
1	0.75	0	0.25	3.52	-0.28%	6.29	0.29%	0.01%
1	0.75	0.25	0	3.58	1.50%	6.39	1.92%	1.71%
1	1	0	0	3.58	1.50%	6.39	1.92%	1.71%
2	0	0	1	3.54	0.29%	6.02	-3.90%	-1.80%
2	0	0.25	0.75	3.54	0.30%	6.02	-3.91%	-1.80%
2	0	0.5	0.5	3.53	0.28%	6.04	-3.66%	-1.69%
2	0	0.75	0.25	3.54	0.32%	6.04	-3.58%	-1.63%
2	0	1	0	3.59	1.75%	6.44	2.75%	2.25%
2	0.25	0	0.75	3.52	-0.28%	6.33	1.05%	0.38%
2	0.25	0.25	0.5	3.52	-0.23%	6.34	1.15%	0.46%
2	0.25	0.5	0.25	3.51	-0.32%	6.32	0.82%	0.25%
2	0.25	0.75	0	3.58	0.66%	6.33	1.68%	1.33%
2	0.5	0	0.5	3.52	-0.28%	6.33	1.05%	0.38%
2	0.5	0.25	0.25	3.51	-0.35%	6.33	1.00%	0.32%
2	0.5	0.5	0	3.58	1.55%	6.42	2.51%	2.03%
2	0.75	0	0.25	3.52	-0.28%	6.33	1.05%	0.38%

Table 6 (continued)

Strategic dispatching metrics				On-demand ride metrics		Food delivery metrics		Avg RIR
MaxBiddingTime	w_urgency	w_wait	w_proximity	Avg_ wait [min]	% Impr. Avg	Avg_ wait [min]	% Impr. Avg	Avg RIR _ Avg
2	0.75	0.25	0	3.58	1.55%	6.42	2.51%	2.03%
2	1	0	0	3.58	1.55%	6.42	2.51%	2.03%

Each row represents a certain parametrization of the strategic dispatching algorithm which includes the duration of the bidding process, *MaxBiddingTime*, as well as the calibration weights, $w_{urgency}$, w_{wait}, $w_{proximity}$. Columns show the average wait time for on-demand rides and food deliveries within that scenario, and how it compares with the scenario without strategic dispatching. The last column represents the average Relative Improvement Ratio (Avg_{RIR}) as defined in Section "Submodels". The row in bold font represents the highest Avg_{RIR}

Performance metrics

See Fig. 14.

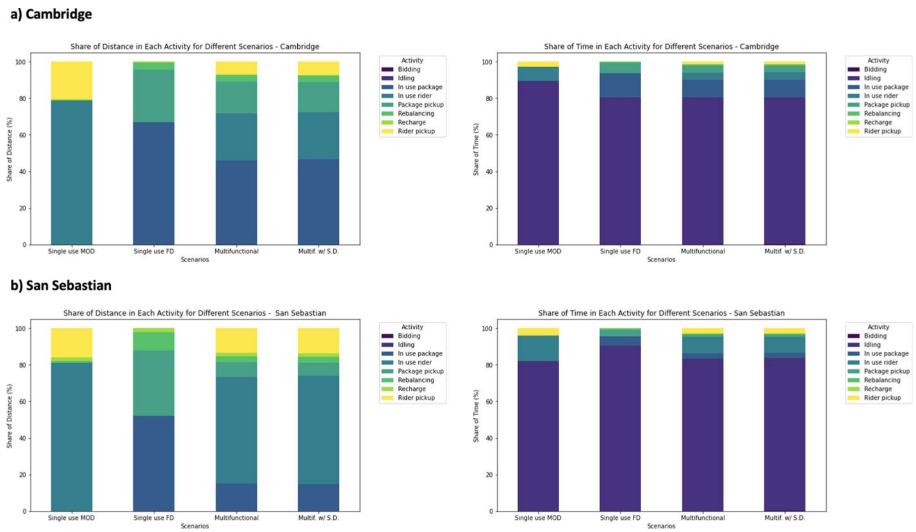

Fig. 14 Performance metrics of the case study of **a** Cambridge and **b** San Sebastian with the calibrated fleet sizes and strategic dispatching algorithm. Figures on the left represent the percentage of the total traveled distance spent on each activity type. Figures on the right, instead, represent the percentage of total time spent on each activity. Each figure shows the results for each of the modeled scenarios: single-use fleets for on-demand rides, single-use fleets for food deliveries, multifunctional service, and multi-functional service with strategic dispatching (left to right)

Acknowledgements We thank the GAMA Platform developer community for their support in this project. We also thank dBizi and Glovo for providing the necessary data for this research. In in particular, we thank Fermin Echarte and Biagio Antonelli from these institutions. Additionally, we thank Andres Rico for his insightful feedback throughout this project.

Author Contributions Conceptualization: Naroa Coretti Sanchez and Kent Larson; Methodology, Software, Formal analysis and investigation, Writing: Naroa Coretti Sanchez; Funding acquisition and Supervision: Kent Larson.

Funding 'Open Access funding provided by the MIT Libraries’
'Open Access funding provided by the MIT Libraries'

Data Availability This research draws upon publicly available data sources, including OpenStreetMap and Bluebikes, as well as proprietary data from dBizi and Glovo. Access to the third-party datasets is contingent upon the policies of these respective companies.

Declarations

Competing interests Authors have no conflicts of interest to declare.

AI-assisted technologies During the preparation of this work, the authors used ChatGPT to improve readability and language. After using this service, the authors reviewed and edited the content as needed and take full responsibility for the content of the publication.

References

Ahuja, K., Chandra, V., Lord, V., Peens, C.: Ordering in The Rapid Evolution of Food Delivery. McKinsey & Company, New York (2021)

Alfeo, A.L., Ferrer, E.C., Carrillo, Y.L., Grignard, A., Pastor, L.A., Sleeper, D.T., Cimino, M.G., Lepri, B., Vaglini, G., Larson, K., *et al.:* Urban swarms: A new approach for autonomous waste management. In: 2019 International Conference on Robotics and Automation (ICRA), pp. 4233–4240 (2019). IEEE

Allam, Z., Bibri, S.E., Chabaud, D., Moreno, C.: The ‘15-minute city’concept can shape a net-zero urban future. Humanities and Social Sciences Communications **9**(1), 1–5 (2022)

Alli, G., Formentin, S., Savaresi, S.M.: On the suitability of epacs in urban use. IFAC Proc. Vol. **43**(18), 277–284 (2010)

Bluebikes: Bluebikes System Data (2023). https://www.bluebikes.com/system-data

Browning, R.C., Baker, E.A., Herron, J.A., Kram, R.: Effects of obesity and sex on the energetic cost and preferred speed of walking. J. Appl. Physiol. **100**(2), 390–398 (2006)

Calabrò, G., Le Pira, M., Giuffrida, N., Inturri, G., Ignaccolo, M., Correia, G.H.D.A.: Fixed-route vs, demand-responsive transport feeder services: an exploratory study using an agent-based model. J. Adv. Transp. (2022). https://doi.org/10.1155/2022/8382754

Chen, T.D., Kockelman, K.M., Hanna, J.P.: Operations of a shared, autonomous, electric vehicle fleet: implications of vehicle & charging infrastructure decisions. Transp. Res. Part A: Policy Practice **94**, 243–254 (2016)

Coretti Sanchez, N., Pastor, L.A., Larson, K.: Can autonomy make bicycle-sharing systems more sustainable? An environmental impact analysis. Transp. Res. Part D: Transp. Environ. **113**, 103489 (2022)

Coretti Sanchez, N., Martinez, I., Pastor, L.A., Larson, K.: On the performance of shared autonomous bicycles: a simulation study. Commun. Transp. Res. **2**, 100066 (2022)

Coretti Sanchez, N., Martinez, I., Pastor, L.A., Larson, K.: On the simulation of shared autonomous micromobility. Commun. Transp. Res. **2**, 100065 (2022)

Coretti-Sanchez, N., Múgica, J., Bowers, Q., Grignard, A., Ferrer, E.C., Alonso, L., Larson, K.: Urban mobility swarms: towards a decentralized autonomous bicycle-sharing system. In: 2023 IEEE 26th International Conference on Intelligent Transportation Systems (ITSC), pp. 2323–2330 (2023). IEEE

Crippa, M., Guizzardi, D., Banja, M., Solazzo, E., Muntean, M., Schaaf, E., Pagani, F., Monforti-Ferrario, F., Olivier, J., Quadrelli, R., et al.: CO2 emis-sions of all world countries-2022 Report, EUR 31182 EN. EUR

Dbizi: Dbizi system data (2023). https://www.dbizi.eus/

De Capitani Da Vimercate, A.: Autonomous robots in food delivery: a simulation study (2019)

Fagnant, D.J., Kockelman, K.M.: The travel and environmental implications of shared autonomous vehicles, using agent-based model scenarios. Transp. Res. Part C: Emerg. Technol. **40**, 1–13 (2014)

Genua Cerviño, A., Coretti Sanchez, N., Wang, E.L., Grignard, A., Larson, K.: Shared lightweight autonomous vehicles for urban food deliveries: a simulation study. Fut. Transp. **4**(2), 634–658 (2024)

Grignard, A., Alonso, L., Taillandier, P., Gaudou, B., Nguyen-Huu, T., Gruel, W., Larson, K.: The impact of new mobility modes on a city: a generic approach using abm. In: Unifying Themes in Complex Systems IX: Proceedings of the Ninth International Conference on Complex Systems 9, pp. 272–280 (2018). Springer

Grignard, A., Taillandier, P., Gaudou, B., Vo, D.A., Huynh, N.Q., Drogoul, A.: Gama 1.6: advancing the art of complex agent-based modeling and simulation. In: International Conference on Principles and Practice of Multi-agent Systems, pp. 117–131 (2013). Springer

Grimm, V., Railsback, S.F., Vincenot, C.E., Berger, U., Gallagher, C., De Angelis, D.L., Edmonds, B., Ge, J., Giske, J., Groeneveld, J., et al.: The odd protocol for describing agent-based and other simulation models: a second update to improve clarity, replication, and structural realism. J. Artif. Soc. Soc. Simul. (2020). https://doi.org/10.18564/jasss.4259

Haas, I., Friedrich, B.: Developing a micro-simulation tool for autonomous connected vehicle platoons used in city logistics. Transp. Res. Procedia **27**, 1203–1210 (2017)

Herrmann, S., Schulte, F., Voß, S.: Increasing acceptance of free-floating car sharing systems using smart relocation strategies: a survey based study of car2go hamburg. In: Computational Logistics: 5th International Conference, ICCL 2014, Valparaiso, Chile, September 24–26, 2014. Proceedings 5, pp. 151–162 (2014). Springer

Huang, F.-H.: Understanding user acceptance of battery swapping service of sustainable transport: an empirical study of a battery swap station for electric scooters, taiwan. Int. J. Sustain. Transp. **14**(4), 294–307 (2020)

International Transport Forum (ITF): Good to Go? Assessing the Environmental Performance of New Mobility Corporate Partnership Board Report. Technical report, International Transport Forum (ITF) (2020)

Jacobs, J.: The Death and Life of Great American Cities. Random House, New York (1961)

Jaramillo, P., Ribeiro, S.K., Newman, P., Dhar, S., Diemuodeke, O.E., Kajino, T., Lee, D.S., Nugroho, S.B., Ou, X., Strömman, A.H., Whitehead, J.: Climate Change 2022: Mitigation of Climate Change. Contribution of Working Group III to the Sixth Assessment Report of the Intergovernmental Panel on Climate Change

Jensen, P., Rouquier, J.-B., Ovtracht, N., Robardet, C.: Characterizing the speed and paths of shared bicycle use in Lyon. Transp. Res. Part D: Transp. Environ. **15**(8), 522–524 (2010)

Kondor, D., Zhang, X., Meghjani, M., Santi, P., Zhao, J., Ratti, C.: Estimating the potential for shared autonomous scooters. IEEE Trans. Intell. Transp. Syst. **23**(5), 4651–4662 (2021)

Li, J., Rombaut, E., Vanhaverbeke, L.: A systematic review of agent-based models for autonomous vehicles in urban mobility and logistics: possibilities for integrated simulation models. Comput. Environ. Urban Syst. **89**, 101686 (2021)

Lin, M.C.-L.: Affordable Autonomous Lightweight Personal Mobility. PhD thesis, Massachusetts Institute of Technology (2021)

Manoeva, D., Salzer, S., Schmidt, S.: An investigation into design engineering aspects of a shared autonomous micro vehicle concept (2022)

Narayanan, S., Chaniotakis, E., Antoniou, C.: Shared autonomous vehicle services: a comprehensive review. Transp. Res. Part C: Emerg. Technol. **111**, 255–293 (2020)

OpenStreetMap: (2023). https://planet.osm.org

Poeting, M., Schaudt, S., Clausen, U.: Simulation of an optimized last-mile parcel delivery network involving delivery robots. In: Interdisciplinary Conference on Production, Logistics and Traffic, pp. 1–19. Springer, Berlin (2019)

Pörtner, H.-O., Roberts, D.C., Poloczanska, E.S., Mintenbeck, K., Tignor, M., Alegría, A., Craig, M., Langsdorf, S., Löschke, S., Möller, V., et al.: Ipcc, 2022: Summary for policymakers (2022)

Romano Alho, A., Sakai, T., Oh, S., Cheng, C., Seshadri, R., Chong, W.H., Hara, Y., Caravias, J., Cheah, L., Ben-Akiva, M.: A simulation-based evaluation of a cargo-hitching service for e-commerce using mobility-on-demand vehicles. Fut. Transp. **1**(3), 639–656 (2021)

Samouh, F., Gluza, V., Djavadian, S., Meshkani, S., Farooq, B.: Multimodal autonomous last-mile delivery system design and application. In: IEEE International Smart Cities Conference (ISC2) (2020), pp. 1–7. IEEE (2020)

Sanchez, N.C., Pastor, L.A., Larson, K.: Autonomous bicycles: a new approach to bicycle-sharing systems. In: 2020 IEEE 23rd International Conference on Intelligent Transportation Systems (ITSC), pp. 1–6. IEEE (2020)

Sayarshad, H.R., Gao, H.O.: Optimizing dynamic switching between fixed and flexible transit services with an idle-vehicle relocation strategy and reductions in emissions. Transp. Res. Part A: Policy Pract. **135**, 198–214 (2020)

Schlenther, T., Martins-Turner, K., Bischoff, J.F., Nagel, K.: Potential of private autonomous vehicles for parcel delivery. Transp. Res. Rec. **2674**(11), 520–531 (2020)

Sheppard, C.J., Bauer, G.S., Gerke, B.F., Greenblatt, J.B., Jenn, A.T., Gopal, A.R.: Joint optimization scheme for the planning and operations of shared autonomous electric vehicle fleets serving mobility on demand. Transp. Res. Rec. **2673**(6), 579–597 (2019)

Taillandier, P., Gaudou, B., Grignard, A., Huynh, Q.-N., Marilleau, N., Caillou, P., Philippon, D., Drogoul, A.: Building, composing and experimenting complex spatial models with the Gama platform. GeoInformatica **23**, 299–322 (2019)

Tech Crunch: Uber is exploring autonomous bikes and scooters. https://tinyurl.com/2dww76er (2019)

Tortoise: Tortoise Shared Micromobility. https://www.tortoise.dev/shared-micromobility (2023)

United Nations Human Settlements Programme (UN-Habitat): Envisioning Future Cities: World Cities Report 2022. United Nations (2022)

US-Foods: New food delivery study shows what consumers crave in a service. US Foods (2019)

Wang, S., Correia, G.H.D.A., Lin, H.X.: Exploring the performance of different on-demand transit services provided by a fleet of shared automated vehicles: an agent-based model. J. Adv. Transp. (2019). https://doi.org/10.1155/2019/7878042

Wooldridge, M.: An Introduction to Multiagent Systems. Wiley, New York (2009)

Zhang, W., Guhathakurta, S.: Residential location choice in the era of shared autonomous vehicles. J. Plan. Educ. Res. **41**(2), 135–148 (2021)

Naroa Coretti Sanchez is a Ph.D. candidate at the MIT Media Lab's City Science research group. Her research focuses on designing future mobility modes and understanding their impacts through modeling and simulations. She aims to enhance the walkability and human-centric nature of future cities by incentivizing increased use of active mobility modes.

Kent Larson is a Professor of the Practice and Director of the MIT Media Lab's City Science research group. With a primary focus on urban innovation, Larson's interdisciplinary work integrates architecture, urban planning, engineering, and computer science to address complex urban challenges. Through his work, Larson contributes to shape the discourse on the future of cities and the role of technology in creating more sustainable, efficient, and livable urban environments.