

The demand potential of shared autonomous vehicles: a large-scale simulation using mobility survey data

ARTICLE HISTORY

Compiled April 3, 2023

ABSTRACT

Shared Autonomous Vehicles (SAV), or robotaxis, are expected to be commercially available within this decade. This new transport mode has the potential to revolutionize travel, offering a level of service comparable to traditional taxis with much lower prices. This may attract travelers currently using other modes, impacting the economic sustainability of public transport as well as car ownership levels. We investigate the scale of this potential demand using a scalable SAV simulation framework. We do not establish a future equilibrium considering the interaction between all users on a detailed road network but establish the potential demand for a large metropolitan area. Travelers can choose between their current mode and the new SAV mode, with an associated fare and waiting times which depend on the real-time demand. For our input data we train a statistical model on a large transport survey from Germany for an urban region, allowing us to generate a large amount of trips with realistic characteristics. We conduct a comprehensive sensitivity analysis to study the effect of several key parameters on the modal shift. A common finding is that SAVs can be attractive to many active mode and public transport users unless regulations are put in place. Our results show also that the SAV fleet constraints can mean that changes in incentives for travelers using currently cars may have significant consequences on the behavior of other travelers. We further calculate key economic indicators for the fleet, which can inform the discussion on what are the fleet size and fare level that operators are likely to choose when maximizing their own profit.

KEYWORDS

Mode shift; simulation; autonomous vehicles; mobility survey; Bayesian network; ride-hailing

1. Introduction

Passenger transport with Shared Autonomous Vehicles (SAV), or robotaxis, are expected to be introduced commercially within this decade. Pilot projects open to the public are already operational in Phoenix, United States, operated by Waymo ([White, 2020](#)), and in Shenzhen, China, operated by AutoX ([Silver, 2021](#)). During the IAA Mobility 2021, Mobileye and Sixt SE announced a plan to deploy robotaxis commercially in 2022 in Munich, Germany ([Mobileye Inc, 2021](#)). This new transport mode has the potential to revolutionize travel ([Sprei, 2018](#)): SAVs can offer the same level of service as traditional taxis or ride-hailing services, but their prices are likely to be much lower, thanks to the absence of the driver. For this reason, demand for SAV services is likely to be significantly higher than for traditional taxis, and some travelers are expected to shift from other modes, such as public transport or private cars.

In this work, we want to investigate the potential scale of this modal shift. We use a utility-based framework to study the potential shift in a population changing from their current mode to using SAVs. We describe the problem as a choice between the currently used mode and such a new transport option assuming that the population might be willing to change their mode in case the SAV provides a cheaper option, in terms of fare and travel time. In particular, our questions are: 1) How might the shift be influenced by the system’s fare and fleet size? 2) How could these system choices in turn influence the profitability of the system, given the shift? 3) How might these choices affect the relative modal shift among modes? The third question is particularly interesting, because different people using different modes have different needs, preferences, and willingness to pay, and their choices influence other passengers in non-linear ways by affecting future waiting times and vehicle locations.

To address these questions, we expand a SAV simulation model to make it more scalable and to incorporate passenger’s discrete choice. We use realistic trip demand data to simulate the choice of passengers between their current mode and the new mode. We generate realistic trip patterns for each mode with a statistical model trained on a large household transport survey of Germany, focusing on a large urban area. Our approach is meant to be a tool to “macroscopically” assess the potential demand. Clearly, our scenarios are hypothetical and not meant to predict modal share in this urban area in a specific year. Rather, we create upper limits as to how many trips might be attracted to SAVs. For a more “microscopic” specific analysis, the approach needs to be extended, and would lead to a data-intensive agent-based model as reviewed in the following section.

We note three key factors here that we do not consider: Firstly, we do not model the detailed congestion impacts on the road network and freed capacities in public transport due to the introduction of SAVs which are likely to cause additional modal shifts among the population. Secondly, additional factors such as “adaptation curves”, and with it trust in SAVs, are not considered. Experiences with other new transport systems, such as high speed rail and car-sharing, have shown that reaching a new market equilibrium will take an extended period of time (if it is ever reached) (Stathopoulos et al., 2017). Axsen and Sovacool (2019) point out the role of users in shaping the SAV technology itself. Thirdly, we limit ourselves to the estimation of modal shift effects, thus we do not consider any induced demand or changes in destinations and activity patterns; nor do we consider potential impacts on urban form (such as people moving from the city center to the suburbs) (Bösch, Becker, Becker, & Axhausen, 2018; Gelauff, Ossokina, & Teulings, 2019). Nevertheless, we suggest this work is useful as a benchmark as to what impacts might be eventually expected without policy interventions. As the results will show, the potential impacts on mass public transport are alarming, if one wants these systems to remain in place.

The work is organized as follows: in Section 2 we summarize the literature on this topic and highlight the original contribution of our work in this context; in Section 3 we present our statistical model trained on survey data for generating realistic trip patterns; in Section 4 we introduce the simulation model used; Section 5 describes how we set up the simulations for the case study; in Section 6 we report the results of our simulations for the case study and discuss their implications; finally in Section 7 we draw our conclusions. Additional information and figures can be found in the attached supplementary materials.

2. Related work and contributions

In this section we summarize the recent literature related to our work, followed by the original contributions of this work in this context.

2.1. Shared autonomous vehicles simulation

There have been by now quite a number of simulation studies on the potential impact of shared (autonomous) vehicles. [Fagnant and Kockelman \(2014\)](#) discuss that autonomous vehicles will quickly increase the market share of car-sharing in line with our hypothesis. They use an agent-based model approach for a grid city to derive waiting times for customers as well as a number of evaluation measures including environmental ones. They highlight the advantages of a reduced vehicle fleet and the impacts of different relocation methods. [Hörl, Ruch, Becker, Frazzoli, and Axhausen \(2019\)](#) simulate SAVs in Zurich with a focus on fleet operational policies and the road network dynamics. They also conclude the large impact of re-positioning strategies on the overall system performance. [Narayanan, Chaniotakis, and Antoniou \(2020\)](#) present a comprehensive literature review regarding modelling SAV scenarios. Important components of SAV modeling include demand, fleet, traffic assignment, vehicle assignment, pricing, charging and rebalancing. Researches often use several data sources for their assumptions with respect to trip request generation.

In many studies (e.g. ([Estandia et al., 2021](#); [Pruckner & Eckhoff, 2020](#))) there is no doubt that future SAV services will be provided by battery electric vehicles leading to the acronym "SAEVs". [Loeb, Kockelman, and Liu \(2018\)](#) use the agent-based simulator MATSim to simulate Austin, Texas, with an SAEV fleet. They discuss charging infrastructure decisions and point out the important impact of different fleet sizes. [Loeb and Kockelman \(2019\)](#) extend the work by considering dynamic ridesharing and a charging station simulation. They suggest that hybrid-electric vehicles might be more suitable in some scenarios. Implications and modeling challenges related to electrification are also reviewed in ([Narayanan et al., 2020](#)). In this case, the complexity of the modeling increases significantly, since both constraints from the mobility sector and from the energy sector have to be taken into account ([Iacobucci, McLellan, & Tezuka, 2018, 2019](#); [Melendez, Das, & Kwon, 2020](#)). We further point out that there are many value added service options for SAEVs in the energy sector. For instance, ([Iacobucci, Donhauser, Schmöcker, & Pruckner, 2021](#)) show that a fleet of SAEVs can provide frequency control reserve, and [Estandia et al. \(2021\)](#) show that overloading and voltage drops can drastically be reduced by coordinating charging under load balancing constraints. Also [Liao, Taiebat, and Xu \(2021\)](#) point out that SAEVs with V2G services have a substantially lower operating cost considering the whole life cycle. While we acknowledge that SAV fleets will likely be electric, in this work we do not consider electrification to better focus on the modal shift.

Rebalancing is one of the major tasks when it comes to Autonomous Mobility-on-Demand (AMoD) systems. Rebalancing problems are often formulated in a station-based context, with stations usually representing areas of the city ([Carron, Seccamonte, Ruch, Frazzoli, & Zeilinger, 2021](#)). Some works use a deterministic approach to minimize waiting times ([Zhang, Rossi, & Pavone, 2016](#)), however these approaches suffer from limited scalability and poor robustness to uncertainty. A faster and more robust queueing-theoretical approach is presented by [Zhang and Pavone \(2016\)](#). [Carron et al. \(2021\)](#) present a discrete-time model of an AMoD system and propose a model

predictive control scheme which can be applied to large-scale cities. Results show that their model predictive control scheme outperforms existing rebalancing schemes for identical fleet sizes. In these scalable approaches the movement of vehicles are considered in aggregate, and a separate assignment algorithm is needed to assign specific vehicles to rebalancing tasks. In terms of transport system impacts, these predictive rebalancing strategies decrease waiting time but significantly increase the empty vehicle-km driven, which in certain cases more than double. [Iacobucci, Bruno, and Schmöcker \(2021\)](#) use a similar aggregate approach to the fluid model in [Zhang and Pavone \(2016\)](#), which we also use in the current work with the addition of relocation clusters to reduce computational complexity.

2.2. SAV demand and modal shift

Future scenarios for SAV simulations are based on travel survey data ([Amanatidis, Langdon, & Clarkson, 2018](#); [Iacobucci, Donhauser, et al., 2021](#); [Iacobucci et al., 2018](#)), land-use and transportation data ([Alam & Habib, 2018](#)), travel records from transit smart cards ([Shen, Zhang, & Zhao, 2018](#)) or taxi usage datasets ([Bauer, Greenblatt, & Gerke, 2018](#); [Iacobucci, Bruno, & Schmöcker, 2021](#)). A range of studies, such as ([Matalqah, Shatanawi, Alatawneh, & Mészáros, 2022](#)) also avoid detailed demand estimation and employ fixed demand conversion rates for origin-destination pairs in order to show the sensitivity to assumptions for traffic conditions. In order to estimate the demand for SAV services many studies assume constant values, use mode choice models or simple growth factor rules applied to these data. For the generation of spatially and temporally random trip requests required as input these data are then combined with the assumption of simple distributions (e.g. Poisson distribution).

To predict modal shift among the generated trips some studies have used or collected additional data. [Mo, Wang, Moody, Shen, and Zhao \(2021\)](#) used a stated preference survey in Singapore to investigate how subjective evaluations and inertia from existing travel modes impact the adoption of AMoD. They incorporate subjective evaluations and inertia into a hybrid choice model. Results show that people with a positive evaluation of ride-hailing and current car drivers are more likely to choose AMoD while users of public transit (PT) are less likely to use AMoD.

A study focusing on the impact of mode choice for long distance travel is presented by ([LaMondia, Fagnant, Qu, Barrett, & Kockelman, 2016](#)). The authors analyze a travel survey and develop a simulation experiment of long-distance model utilizing a long-distance trip generation model and a long-distance mode choice model. The simulation experiments reveal interesting insights in the potential shifts in mode choices across different trip distances and purposes.

[Heilig, Hilgert, Mallig, Kagerbauer, and Vortisch \(2017\)](#) analyze the potentials of SAVs in a changing private transportation system by utilizing an agent-based microscopic travel demand model called *mobiTopp* to simulate the mode choice behavior in a case study for the Stuttgart region. In their main scenario they assume a world without private cars and the presence of large SAV services instead. Findings show that not all SAV services can substitute all trips previously made by private cars, meaning the modal share of bicycling, walking and PT is increasing as well.

[Mo, Cao, Zhang, Shen, and Zhao \(2021\)](#) investigate the competition between SAVs and PT. The authors study several scenarios to examine the competitive perspective of SAV and PT operators and find that the competition can result in higher profits and higher system efficiency for both operators. From passengers perspective the com-

petition can reduce the travel time and travel cost when considering the value of time. Results show also that AV operations should be regulated with regard to number of licenses, operation time and service areas.

[Gao et al. \(2021\)](#) analyze the commuting travel shift behavior by investigating inertia effects of past behavior. Thereby the authors use an empirical, mixed dataset of revealed preference and stated preference surveys. The authors show that there are huge interactions between inertia effects and level of service variables. They conclude that inertia effects are also an important influencing factor in traffic assignment models.

[Golbabaei, Yigitcanlar, and Bunker \(2021\)](#) provide a review of the literature on the impacts of SAVs on urban mobility, including factors influencing mode choice towards SAVs. They highlight how many studies found that lower costs and the higher value of time gained compared to private cars may drive adoption of SAVs.

[Tian, Feng, Timmermans, and Yao \(2021\)](#) utilize a stated choice experiment to gain insight into individuals' propensity for purchasing an autonomous vehicle or using shared (autonomous) cars. The authors show that substantial heterogeneity exists among individuals in their preferences, generally people prefer to keep their current car or buy an AV rather than to use shared cars. From a demographic point of view, younger people are more open to accept new mobility options.

[Hörl, Becker, and Axhausen \(2021\)](#) utilize the MATSim agent-based simulation platform to model a system of automated taxis in Zurich. The authors conclude that an unregulated SAV system could capture a large share of transport demand not only from private cars, but also public transport and active modes, leading to increased congestion, noise, and accidents. They suggest several policy measures to limit the negative impacts, such as operational area restrictions, incentives to share vehicles and using SAVs for multi-modal trips with public transport. The approach used allows for detailed simulation of the impact of the new mode on the urban transport in the city. However, it requires significant amount of data which is often not available. The MATSim mode choice model based on utility scoring has also been used in some of the aforementioned studies as well as, for example, the study of [Vosooghi, Puchinger, Jankovic, and Vouillon \(2019\)](#) who simulate SAVs being used for individual trips as well as ride-shared trips.

In order to support planners and help them to adjust SAV policies, [Etminani-Ghasrodashti, Ketankumar Patel, Kermanshachi, Rosenberger, and Foss \(2022\)](#) implementing SAVs as a new mobility mode in a simulation model, where also usual traditional transportation mode are modeled. Results show that people with greater access to a private vehicle were more interested in using private vehicles and less interested to use ridesharing alternatives. Other studies, such as, [Huang, Kockelman, Garikapati, Zhu, and Young \(2021\)](#), focus on the potential shift when SAVs are providing access and egress services to a transit station. They use the SUMO simulation framework and its inbuilt mode choice model to simulate the potential demand shift from persons currently walking, or driving to the stations as well as the potential for persons currently driving all the way.

[Stoiber, Schubert, Hoerler, and Burger \(2019\)](#) conducted an online choice experiment with over 700 participants from Switzerland to study future mode choices under the assumption of full-market penetration of autonomous vehicles. The experiment tested the influence of short-term and long-term decision instruments to encourage the adaption of SAVs. Results showed that around 60% of Swiss respondents prefer pooled autonomous vehicles over private AV.

Similar to [Stoiber et al. \(2019\)](#), [Haboucha, Ishaq, and Shiftan \(2017\)](#) also utilized a

stated preference questionnaire distributed to over 700 participants living across Israel and North America. The vehicle choice model provide three options including continue to commute using a regular car, buy and shift to a private-owned autonomous vehicle or shift to SAV. Results show, that only 75% of individuals are interested in using SAV even if the SAV service were to be completely free.

Eom, Lee, Ko, and Lee (2022) analyze a mode choice model based on socioeconomic characteristics of individual travelers in a smart city context of Korean city of Sejong. In this study, the authors consider different transportation modes such as private cars, bus rapid transit, carsharing and ridesharing. Results show that around two-thirds of respondents in the survey preferred to spend less than 1.25 USD and travel less than 35 minutes on bus rapid transit trips.

A related line of research is that of the shift towards shared micromobility. E-scooter sharing in particular is becoming increasingly popular. Initial work already exists that focuses on the factors for shifting to e-scooter sharing, particularly in the micromobility sector. For example, Cao, Zhang, Chua, Yu, and Zhao (2021) focuses on e-scooter sharing to replicate short-tips to public transport in Singapore. Another specific question on the shifting behavior of users to e-scooters is posed by Abouelela, Al Haddad, and Antoniou (2021). In their work, the authors show that there are several factors influencing the shifting behavior. For this purpose, they conducted a stated preference survey among young adults to estimate a choice model between carsharing and scooter-sharing. This model is applied to evaluate the switching behavior using a real carsharing dataset. Results show that up to 23% of carsharing demand can be met by scooters. The current results on mode shift in the field of e-scooters show the enormous potential of new, smart mobility services. *Therefore, in this study we show potential upper limits as to how many persons and which trips might be attracted towards SAVs could look like based on monetary and temporal criteria.*

2.3. Contributions

Given this literature review, we suggest the main distinguishing points of this work compared to earlier contributions are as follows:

- (1) We present a novel method to generate trips for a whole city from limited survey data by using Bayesian networks. While augmentation methods have been presented before, we suggest a strength of the proposed approach is its ability to consider a range of correlated trip attributes such as distance, purpose, starting time and modality. The approach is able to repeatedly generate millions of trips in a metropolitan area considering these interrelationships. *Because we create “trips” and not “tours”, we do not claim that our method is more accurate in reflecting all the daily constraints that determine mode choice decisions; but we suggest that our approach is an efficient, fast way to provide an estimate as to what trips in a city might be attracted to SAVs.*
- (2) Accordingly, we are able to investigate the effects of shared autonomous vehicles on an urban area using realistic trip demand and current mode choice patterns for a whole week. This allows us to frame our study as a choice between the currently chosen mode alternative and the assumed newly available SAVs. We also include as original modes not only cars and public transport, but also walking and cycling. As the literature review showed, few studies have done so and instead often focused on modal shift from a particular mode.
- (3) Furthermore, whereas other studies have focused either on long-distance travel




Figure 1. Flowchart of the work. Blue diamonds are the major external input sources, red rectangles the major models, and black ellipses are the data outputs. The black rectangle represent the input parameters for the simulation

or the urban core only, we include the surrounding sub-urban region. We suggest this metropolitan area perspective is important as the attractiveness for SAVs will depend on the level of public transport offered in the different parts of the modeled area.

- (4) Our SAV demand allocation and predictive relocation model is based on previous work. In those studies the focus has, however, not been on potential modal shift scenarios. Here we perform a comprehensive sensitivity analysis to study the effect of several important parameters on the mode shift. In particular we did not find in previous work a distinction of short-term and long-term private car costs in the discussion as to how SAV adaption patterns might differ in medium and long-term future scenarios.

To achieve these contributions we developed the framework shown in Figure 1. The different elements of this framework are explained in the following Sections 3 and 4. We note that we could also reduce the scope of our paper to the simulation model and estimated model shift only. In that case, to obtain our modal share estimates for a whole metropolitan region we require an external trip generation model that covers the whole area and is detailed enough to include the current mode choice and an estimate of the cost of the currently used mode. As reviewed in [Narayanan et al. \(2020\)](#) such “external demand generation models” for existing SAV simulations are based on a range of data including travel survey data (as in this work) or estimates from land-use and transport models. We argue, however, that none of the existing approaches reported in the literature is very suitable for our purposes, since the region and trip numbers covered in our approach are large and we cover potential modal share from all currently used modes. We further want a flexible, fast trip generation model that allows us to run several scenarios. To illustrate this, consider the SAV simulation of [Hörl et al. \(2019\)](#). Input for their MATSim model is a prior calibrated demand model

based on a Swiss National survey, comparable to our survey data. For our, compared to MATSim in several aspects less detailed, simulation we want instead a fast trip generation model that is able to create only the key demand characteristics that we use in our simulation with millions of agents. The trip generation model should further be flexible to create the future demand for a range of scenarios.

3. Trip generation

In this section we introduce our statistical model to generate any desired number of trips based on a person trip survey. We expand the formulation in [Iacobucci, Donhauser, et al. \(2021\)](#) to include mode split and an improved clustering approach. Since the implementation of the model is dependent on the data available, we firstly present our dataset in Section 3.1, then we describe the model in detail in Sections 3.2–3.7. Finally, we report the validation of the generated data in Section 3.8. An outline of our methodology is in Figure 2.

3.1. Data set

We base our analysis on a typical person-trip survey conducted in many countries. For our case study we take the largest transport survey in Germany, the *Mobilität in Deutschland* (MiD), or *Mobility in Germany* study ([Follmer & Gruschwitz, 2019](#)). The MiD dataset contains almost 1 million trips of 316,361 persons during the reports observation period, starting in May 2016 until September 2017. Although the survey covers the whole of Germany, the availability of data varies depending on the geographic area. In particular, the density of trips is much higher close to large urban centers compared to rural areas. For this reason, in this work we focus on a city and its surrounding suburbs.

Assuming representativeness, we can use the spatial information contained in the dataset as ground truth for our model. For our purposes of simulating an SAV fleet, it is, however, important to scale up the sample to represent all trips.

We chose to focus on the metropolitan area of Munich, which is one of the best covered areas in the survey with 50,000 survey entries. We consider a square area within 17km to the city center (33x33 km), which includes the city’s highway ring. Figures S4 and S5 in the supplementary material report the number of trips by mode and purpose, respectively. Note that the average number of samples per grid cell is 46, which is quite low. This is why the augmentation approach presented here is necessary. Although in this work we focus on its implementation for Munich, the approach presented here is general and can be applied to any city or region, assuming that sufficient data are available.

The origin and destination of the trips are given in a grid structure with square grid cells with a resolution of 1km. Not every trip has complete information. To resolve this problem with missing data without introducing biases in the mode patterns we perform a “pre-augmentation” step, described in Section S-I.1.

3.2. Trip generation model overview

The model is designed in such a way that it is able to capture central properties of person trip datasets, such as trip purposes, modalities, and start-time, with dependencies




Figure 2. Pipeline for data augmentation. The local data is clustered with *k-Means* and a Bayesian Network is fitted according to the data, with origin and destination as a cluster-id. Both variables are then mapped into continuous space with cluster wise *KDE*, time is made continuous using GMMs.

to the origin and destination location using sparse data.

Due to the high spatial resolution of the survey, there are many origin and destination cell pairs where no trips are present in the dataset. When further conditioned by hour, the triplet (origin, destination, hour) is never seen for most cases. Thus, it is clear that simple statistical counting methods can not be used to fill in the detailed, sparse matrix. While one could "simply" fix this by switching to a coarse representation, here we instead propose a clustering based method which groups the space into regions which are smaller if enough data is available (e.g. in the city center) and bigger if the data is more sparse. In this way the model can adapt to the spatial availability of data.

The core unit of the approach is a Bayesian Network (BN), which captures probabilistic dependencies of the MiD data, as in OD pairs, trip purpose and time. Due to the low availability of data, the BN does not work with the exact spatial information but is operating on cluster level instead. Those clusters can be seen as differently behaving regions, e.g. the destination region depends on the purpose.

To get a continuous time/space domain (instead of discrete hours and cluster positions) from the model, a refinement step is performed. The continuous time is sampled by a Gaussian Mixture Model (GMM) fitted to the data, while the discrete clusters are made continuous using Kernel Density Estimation (KDE).

From the continuous origin and destination pairs the (Euclidean) trip-distance is calculated, which we found to be highly correlated to the trip-modality (see also Figure S6 in the supplementary material). This statistical dependency is then used to guess a likely modality for the trip, which we have ignored in the core-BN to save an additional data split.¹

We implemented the model with Python. For details of the implementation, refer to Section S-I.3 in the supplementary material. For brevity, in the following sections we present the tool-chain on a more abstract level. In the implementation, we essentially have two phases: 1) The learning phase, consisting of clustering, Bayesian probability estimation, KDE and GMM fitting; 2) The sampling phase, in which agents are generated according to the BN, and sampling steps from the fitted models are utilized where required.

¹Each data split will reduce the number of trips from origin A to destination B. A simple example: The trips from A to B on Mondays only are less than the trips for all week days.




Figure 3. Location of the cluster centers and the assigned regions for each cluster. In the city center the clusters are tighter because the data availability is better.

3.3. Clustering

The clustering is using the k-means algorithm with location data (coordinates of origin and destination), where multiplicities/duplicates are allowed. By training with multiplicities (i.e. multiple trips with the same origin/destination), the clusters will get tighter around cells with high data availability. The k-means algorithm expects the number of desired clusters N_C as input. Although clusters have an important role, we found that their exact number does not have a significant influence on the outcome. A possible explanation for this is that clusters act as behavioral regions and as tool to get sufficiently large data chunks, but their exact location does not matter too much. From experimental results we found that $N_C = 30$ is suitable for our scenario (see Figure 3).

3.4. Bayesian network

The Bayesian network (BN) presented here is a probabilistic directed acyclic graph, capturing key dependencies between MiD attributes organized in different levels of dependencies. An outline of the BN is shown in Figure 4. The presented structure was not motivated only by statistical considerations, but also from a practical perspective. In particular, the small amount of data available forces us to drop statistical dependencies, otherwise the data would become too sparse which in turn would cause overfitting. For example, in our model the trip start time does not condition any other feature. This is also why the core-BN is built on clusters.²

3.4.1. Core Bayesian network

The top two layers – outer influence and core network – are building the core of the BN. In this part the main trip characteristics are defined. Except for the start-time,

²There are N^2 possible origin-destination pairs for N nodes. While choosing $N^2 := N_C^2 = 30^2 = 900$ seems appropriate for roughly 50000 trips, choosing $N^2 := N_{cells}^2 = 1000^2 = 10^6$ is much more than the amount of data and thus highly unsuited.




Figure 4. Outline of the full Bayesian network used in this work. The network has a hierarchical top-down approach: starting with hard facts/outer influences, to core dependency calculation, to refinement of sampled values, to auxiliary calculations and randomly deciding further attributes using knowledge from the upper layers.

every node here is categorical (e.g. 7 different days). The first layer is representing outer influence on the model. For every possible outer influence (e.g. day=Monday), the data is split and the remaining features are calculated on those splits. This is useful to view the data under different aspects (like different days), but can not be utilized extensively, because the data would be split up to much becoming too sparse even with large surveys. The second layer captures suitable dependencies between origin (cluster), destination (cluster), trip-purpose, and start-time. Obviously there is a dependency between origin and destination, the fact that origin conditions destination is by choice. Note that a bidirectional dependency does not fit the BN-requirement. As mentioned in Section 3.1, the origin-destination dependency is quite heavy and conditioning it further by time would split the data too much, even though it might be more realistic. We found that the discrete trip-purpose (e.g. work) is a suitable alternative, because it correlates with the time and origin/destination. So we condition the destination and the time by the purpose. We found that a custom combination of purpose fields in the MiD dataset produced the best results. One strength of this model setup is that the transition probabilities between the clusters for generated agents are the same probabilities as seen in the training data.

3.4.2. Additional Bayesian network steps

The core network is extended by two further layers, one for spatial refinement and one for additional calculations and further probabilistic guessing. The refinement step is essentially the KDE fitting, explained in section S-I.2.2. Next to the refinement some auxiliary calculations are performed: The Euclidean distance is calculated as well as the origin and destination cell for comparison purposes. The lower level of the BN is suited to sample the modality for the trip. In Figure S6 the dependencies between modality and trip-length and between modality and purpose is shown. We use this to

get a suited modality out of the network without the need to further split the data in the top section.

3.5. Continuous space/time domain

The BN structure described in Section 3.4 includes nodes with a continuous domain (one dimensional time, two spatial dimensions). We use GMMs to fit a continuous time distribution, and we generate continuous space coordinates using KDE. The details of the implementation are reported in Section S-I.2.

3.6. Sampling

For sampling we proceed as follows: the already sampled attributes, e.g. week-day=Monday, purpose=Work, origin-cluster=3 and destination-cluster=7, are used to choose the respective fitted distribution for start-time and the cluster-wise exact origin/destination. From those distributions a sample is drawn randomly.

3.7. Calculating travel times

To provide more realistic agent travel times and distances, the Open Source Routing Machine (OSRM) is used. OSRM is a stand-alone application providing an efficient routing API based on OpenStreetMap data. OSRM is not only able to provide street map based travel times and distances for cars, but also for pedestrians and bicycles. The MiD dataset does contain train and public transport trips, which may also be sampled as agents by the augmentation algorithm. Since OSRM does not provide railway routing, no OSRM based travel time or distance can be assigned to those trips, and an alternative approach is used in these cases, as described in Section 5.3. In the future, the framework could be extended to use data from OpenRailwayMap (<https://www.openrailwaymap.org/>) to fill this missing attribute.

3.8. Validation

For the discussion of the quality of the data augmentation it is firstly important to note that all of the presented results rely on the MiD data as ground truth and do not investigate whether the MiD data itself is representative.

One of the major criteria for the validity of the augmentation is that the spatial distribution of the agents follows the shape of the city and does not blur too much. Figure S2 clearly shows that the density captures the shape of Munich and that the sampled agents are matching the distribution. By comparing Figure S2 to Figure S1, it can be seen how the input data and the generated agents are related: the overall distribution is similar, but the sampled agents are more spread to higher diffusion, which is a wanted effect up to a certain degree. Figures 5–7 show a comparison of the trip time, distance, and modes for the original data and our generated data. The very unnatural shape of the MiD data in Figure 6 suggests that the trip length is mostly a rounded estimation rather than a perfect measure. By comparing the density estimation lines the data for both shows high similarity, which is an important aspect for the overall validity of the presented model and important for the modality guessing.




Figure 5. Comparison between original data and the samples agents. The bars show the share on the time distribution per hour, the two lines show a fitted KDE. The match of the lines is arguably better than the match of the bars. A possible explanation for this is that the real time distribution is not entirely a mixture of Gaussian distributions. Especially the quick rise of traffic during the morning hours seems to violate that constraint.

4. Transport simulation model

We develop a SAV simulation model partly based on the model described in [Iacobucci, Bruno, and Schmöcker \(2021\)](#). The original integrated model includes a simulation model for the movements of vehicles, a vehicle relocation optimization, a charging optimization module, and algorithms for vehicle assignment to passenger requests. The original model also allows simulation of electric vehicles with a specific range, as in [Iacobucci, Pruckner, and Schmöcker \(2022\)](#). However, for simplicity and to better focus on the modal shift, we do not consider range limitation in this work. We extend this model in two main ways: 1) we add a simulation of passengers’ mode choice; 2) we introduce a clustering system for relocation actions. In the remainder of this section, we present the details of this extension. We implemented this model in Matlab.

4.1. Simulation model overview

Our simulation model is based on discretized time and space. The system service area is represented by N nodes, and V vehicles move between nodes to satisfy passenger requests. We assume that passenger requests start and end at the centroid of each node. We model travelers as a sequence of K elements indexed by k . Each traveler k in our simulation is associated with a tuple (t_k, o_k, d_k, u_k) representing the trip time, the origin node, the destination node, and the utility of the alternative (original) mode. Each traveler has a choice to use the SAV system: in that case, a new request is created from its origin to its destination, as described in the next section. The vehicle status is modeled by three time-varying parameters: current vehicle destination $a_v(t)$, current distance to destination $h_v(t)$, and current vehicle status $s_v(t)$. Vehicle status can be one of the three: 1) idle; 2) moving towards or with a passenger; 3) relocating empty for predictive relocation. The notation used here is summarized in Table 1.




Figure 6. Comparison between the trip lengths given in the MiD and the generated agents.




Figure 7. Comparison of the modalities of 20000 sampled agents to the training data. The matching is quite accurate in most cases, even though the modality depends on the indirectly calculated distance.

4.2. Vehicle assignment and mode choice

In order to obtain a SAV demand estimate, we need to know the expected modal shift from other modes to this new ride-hailing mode. In general, the lower the cost of a new mode compared to the current choice, the higher the probability of a traveler to shift to the new mode. In order to simulate the modal shift from an existing mode to the SAV system, we introduce a simulation of passenger’s mode choice with a discrete choice model. We assume that each traveler k is able to obtain accurate information on the waiting time and price of the SAV trip from its origin to its destination at the time of mode choice decision. The traveler then chooses to stick either with their original mode or with a SAV according to a probability that depends on the expected waiting time, travel time, and price of the SAV trip, and its originally chosen mode. We use a binary logit model to simulate the choice of passengers. According to this model, the probability of traveler k choosing SAV over the original mode is given by:

Table 1. Notation

indexes and sets	
v	vehicle
t	time step
k	travelers
model variables	
$a_v(t)$	destination of vehicle v at t
$h_v(t)$	distance to destination of vehicle v at t
$s_v(t)$	status of vehicle v at t
p_k	probability of traveler k choosing SAV
w_k	waiting time offered to traveler k
x_k	binary variable indicating whether traveler k chooses SAV
model parameters	
V	number of vehicles
z	fare of SAV trip (€/km)
t_k	time of travel for traveler k
o_k, d_k	origin and destination nodes for traveler k
u_k	utility of original mode for traveler k
$d_{i,j}$	distance (km) between nodes for SAV
$\tau_{i,j}(t)$	travel time (minutes) between nodes at time t for SAV
VOT	value of time (€/min)

$$p_k = \frac{\exp(u_k^{SAV})}{\exp(u_k^{SAV}) + \exp(u_k)} \quad (1)$$

$$u_k^{SAV} = -\max(z_{min}, z d_{o_k d_k}) - (\tau_{o_k d_k}(t) + w_k) \text{VOT} \quad (2)$$

where z , d , and τ are the price per km, the distance, and the travel time of the trip by SAV; u_k^{SAV} is the utility of the SAV trip for traveler k ; w_k is the waiting time that the SAV system offers to traveler k ; VOT is the value of time. The max term ensures that passengers always pay at least the minimum price for a trip. We did not include the possible utility of traveling on a SAV compared to other modes, due to high uncertainty over this value. However, we can expect that travelers shifting from a private vehicle, such as a car, may prefer the SAV also because the time spent traveling is more valuable compared to time spent driving. The results are therefore somewhat conservative from this point of view. On the other hand, some traveler may have a negative utility associated with SAVs due to concerns over safety or unfamiliarity with new technologies. Quantifying this is also very uncertain and we decided to omit this term as well, since it is also in the opposite direction of the previous.

We assume that each traveler enters the system in order, and the system gives priority in the order of arrival on a first come-first served basis. This is similar to the way current ride-hailing systems generally work (Correia & van Arem, 2016). The algorithm for waiting time calculation, mode choice, and trip assignment is reported in Algorithm 1.

Algorithm 1 Assignment of vehicles to trip requests at time step t

```
1: for each traveler  $k : t_k = t$ , in order of  $k$  do
2:   for each  $v \in \mathcal{V}$  do
3:      $w_k^v = h_v + \tau_{a_v, o_k}(t)$ 
4:   end for
5:   select vehicle  $v'$  with lowest  $w_k^{v'}$ 
6:    $w_k = w_k^{v'}$ 
7:   let  $x_k$  be a random variable with  $Pr(x_k = 1) = p_k$  calculated in (1)
8:   if  $x_k = 1$  then
9:      $h_{v'} = h_{v'} + \tau_{a_{v'}, d_k}(t)$ 
10:     $a_{v'} = d_k$ 
11:   end if
12: end for
```

4.3. Predictive relocation clustering

Predictive relocation moves vehicles closer to the expected location of future requests to minimize waiting times. The original model in [Iacobucci, Bruno, and Schmöcker \(2021\)](#) assumed predictive relocation to occur between each node. This limits the spatial resolution due to increased complexity of relocation operations with increasing number of nodes. To solve this problem, we perform a clustering of the nodes, and we choose the node closest to the center of each cluster as the destination of relocation actions for that cluster. Note that these clusters are different from the clusters used for trip generation in Section 3.3. By performing relocation between clusters of nodes, the spatial resolution can be increased while keeping the computational time low. The imbalance calculation is performed over departures and arrivals in all nodes belonging to a cluster, and vehicles are moved between their current node and a cluster center (relocation destination node).

We base our predictive relocation on the expected number of passengers for each cluster calculated with the same mode choice logic described in Section 4.2 from the original modal split.

4.4. Case study: Munich scenario

As explained in Section 3.1, we consider a square area of 33 by 33km centered in the city of Munich. This area includes 1089 square cells with sides of 1 km, which we take as nodes for the simulations.

We use an exogenous travel time matrix between each node, which changes during the day depending on average traffic conditions on the roads. First, we calculated the distance and travel time without traffic between each node by using the OSRM as described in Section 3.7. This gives us distance and travel time for a total of about 1.2 million OD-pairs in our area. We then cluster our nodes into 100 areas, and we use Google Maps Distances API to estimate car travel times (considering traffic congestion) between the center of each area at three times during the day: 8:00, 12:00, 17:00, which are the morning peak, day off-peak, and evening peak. We can then calculate traffic time multipliers for each area-OD (10,000 in total) and time as the ratio between the travel time in traffic and the travel time without traffic. We calculate the multiplier for each minute of the day as linear interpolation between each time, assuming that there is no traffic congestion between 21:00 and 6:00 (night time). We then multiply




Figure 8. Trip distance by mode (%) for the trips during the week.

the travel time between each OD node pair with the multiplier corresponding to the areas belonging to the origin and destination node, thus giving us the estimated travel time for each OD-pair in each minute during the day. We assume that our fleet does not affect traffic. This is a reasonable assumption, since in most cases considered here the fleet is relatively small compared to the total number of vehicles on the roads in Munich.

We use 100 relocation clusters. We found these clusters with the k-medoids method (Schubert & Rousseeuw, 2019) based on the average travel distance between nodes. This method is superior to the more commonly used k-means clustering for this application, since it can use an arbitrary distance matrix (in our case, travel time) rather than the Euclidean distance. Moreover, the resulting centers are chosen among the data points, thus also giving us the relocation cluster centers.

5. Costs of current mode

To estimate the modal shift from existing modes to the new SAV mode, it is important to know the willingness to pay of passengers, and in particular the current expenditure for travel. In other words, we need the utility for the alternative mode u_k described in Section 4.2. In this section, we detail the methods used to estimate current costs for each mode that travelers are paying.

5.1. Overview

We run our simulation model with a mode choice simulation based on the logit model described in Section 4. We assume that passengers have two alternatives: their current mode, and the SAV mode. In this section we discuss the assumed costs for each current mode. [Before doing so, we re-iterate and expand our discussion on the limitations of the approach which mean we only estimate an upper bound of current trips changing to SAV. For one, we assume that unchosen alternatives before introduction of SAV remain unattractive. Refining this assumption would require us to abandon our scalable approach.](#) Further, we assume people decide the mode for each trip independently of subsequent trips. This is often not the case. For example once an outward trip is made by car or bicycle then the mode of the return trip is generally fixed. This limitation is a result of creating a zonal mode specific OD matrix and not reverting to fully agent-based modeling and considering individuals and their tours. However,

thanks to this approach we can model very large number of trips which would otherwise be impossible with the data available. We also note that most of the effects due to this limitation tend to cancel out over the large number of trips considered. For example, for a Home-Work-Home tour, the probability of shifting for the outbound trip is similar to the inbound trip, since both share the same OD zones with similar characteristics. This limitation becomes more significant with less symmetric demand and more complex tours.

5.2. Active mode and car user costs

The total mode cost will be the sum of a monetary cost (such as tickets for public transport) and a cost related to travel time, which is priced at a specified value of time (VOT). We remind that in our work we consider four modes: walking, cycling, private car, and public transport. In this section we focus on the first three modes.

For the travel time costs, we use the travel time for walking, cycling, and private cars found with OSRM and the Google Maps API as described in Section 4.4.

For car users we also add an average time for finding a parking place. The average time spent searching for parking in Munich is 9 minutes for on-street parking and 6 minutes for off-street parking (INRIX, 2017). We use the average between the two, 7.5 minutes, for all trips by car. Note that this is still an underestimation of the true parking costs, since due to limitation of our data we do not consider the monetary cost of parking, which can be quite expensive in Munich (Agora Verkehrswende, 2019). Short and long term scenarios differ only in terms of private car costs, while costs for all other modes are the same.




Figure 9. Map of rail transport in Munich overlaid with our grid

For the monetary costs, we assume they are zero for walking and cycling. For the monetary cost of private cars we consider two cases:

- (1) **Marginal costs:** This considers only the cost of the trip, which is mainly the fuel cost, since the cost of the car is a sunk cost. We consider a fuel consumption of 5 l/100km and gasoline price of 1.5 €/l.
- (2) **Full costs:** In this case we use the total cost of ownership (TCO), which includes the marginal cost as well as the purchasing price, taxes, etc., divided by the average km driven. We use an average TCO of €600 per month with an average of 10,000 km per year (Letmathe & Soares, 2017), which is equivalent to €0.72/km. This is the average TCO in Germany for a medium size car such as

the Volkswagen Golf (Letmathe & Soares, 2017).

5.3. Public transport costs

We do not have information about which public transport route the passenger used, and it would be impractical to recreate in detail the best route for each of the over 30 million trips. For this reason, we reconstruct an approximate route by considering only the two main public transport modes in Munich: train and bus. For the train mode, we use the metro trains (U-Bahn) and rapid urban transit (S-Bahn) network of Munich and associate each station to a node in our model (Figure 9). For the bus network, we assume that buses travel the same distance as cars between nodes. The average speed of buses in Munich is 17.9 km/h as reported by the Munich transport company MVG ([Münchner Verkehrsgesellschaft mbH \(MVG\)](#), n.d.). In this work we use a more conservative average speed of 15 km/h to account for transfers and traffic during peak hours, which is when most passengers use buses. Moreover, most trips happen in more central areas where average speed is generally lower. We assume that any node can be reached by any other node by bus in this way. This is reasonable given the dense bus network of Munich. We then calculate the travel time for each trip as the best combination of bus and train from its origin to its destination. We also add an average of 5 minutes for each change and 10 minutes to each trip to account for the average waiting time and walking distance to the station or bus stop. This approach is conservative, since we assume a low waiting time and a high frequency throughout the day.

We also associate a fare for each trip, calculated according to the zonal system in Munich. Our simulation area includes nodes located up to the fourth fare zone. We assume that passengers use the Stripe Ticket system, where passengers buy a ticket with 10 ‘stripes’ each worth €1.46, and one or more stripes can be used to travel on the system, depending on the number of zones traversed in the trip. The lowest fare of one stripe is paid for short trips, which are defined depending on number of stops and mode and independent of the zone ([Münchner Verkehrs- und Tarifverbund \(MVG\)](#), n.d.). We estimate which trips will be a “short trip” through our public transport shortest path calculations.

6. Results

We test the model on the generated trips and run all the simulations over a week, with an average of 4.6 million trips per day (32 million trips in total).

6.1. Scenario analysis

We analyze several scenarios to understand the possible impact on modal shift and the consequences for the fleet operator. We consider:

- (1) four fleet sizes (2,500, 5,000, 10,000, 20,000)
- (2) three SAV pricing levels (0.5, 1, 2 €/km)
- (3) two car cost scenarios (marginal costs and full costs, defined in Section 5.2).

The performance of our model in terms of computational time is reported in Figure S7.

Fleet size is one of the most important parameters influencing modal shift and economic sustainability of the system. To test its influence we run our model with different fleet sizes ranging from 2,500 vehicles (which corresponds to one vehicle for every 1,830 trips per day made by the population) to 20,000 vehicles (one vehicle for every 230 trips per). Note that a fleet of 20,000 vehicles is probably unrealistic giving our assumption of no induced congestion from the SAV system. This large fleet is likely to have a significant effect on congestion, even when replacing a large proportion of private cars currently on the road. Despite this, we include this scenario for comparison with other fleet sizes, but keeping in mind that the modal shift from non-car based modes might be overestimated.

Another key parameter is the price per km. The lowest price we consider, 0.5 €/km, is similar to the equivalent cost per km of car-sharing systems,³ thus it can be considered a lower bound for the economic feasibility of these systems. On the other hand, the highest price considered of 2 €/km is closer to the price of taxis and ride-hailing systems, thus can be considered as a suitable upper bound: Operation costs of SAV systems are likely to be significantly lower (and thus could be priced lower) thanks to the lack of drivers. The size of the fleet and the price per km in a real system will depend on economic considerations from the fleet operators. For this reason, in our results we also report some basic economic indicators, which should suggest what are the most likely combinations of fleet size and pricing to be implemented.

Table 2 summarize the results for the four fleet sizes, three pricing strategies, and the two cases with the ‘marginal’ and ‘full’ car cost cases defined in Section 5.2. Figure 10 shows the average modal shift by mode for the case with a price of 1 €/km, and Figure 11 reports some economic indicators for the same case. The total modal shift (and therefore potential demand of SAV from current trips) varies between 3.7% and 27.5%, with larger fleets and lower prices able to attract more passengers. However, larger fleets are also less efficient from the point of view of the fleet operator, with less vehicle utilization (less vehicle-trips per day, VT and less vehicle-km traveled per day, VKT) and less revenue per vehicle (RV). The share of empty vehicle-km traveled (eVKT) varies substantially between cases, ranging from 7% to over 17%. This mostly depend on the type of trips that the system is able to attract. More asymmetric trips lead to higher eVKT. For the system operator, the solution with the highest profit per vehicle is reached with a fleet size of 2,500 vehicles and a 2€/km fare. The solution with the highest total profit is the one with 20,000 vehicles and 1€/km. For details, see Figure S8.

In all cases, a large portion of walking trips shift to SAV. This is partly a result of our spatial resolution: our model only captures the longest walking trips, since we only consider trips between different nodes (interzonal trips only) and each node represents an area of 1 km squared (see also Figure 8). Shorter walking trips are less likely to shift to SAV due to the non-zero waiting time and minimum price. (Note that we therefore call the corresponding column in Table 2 “Interzonal foot trips”).

The shift from cars is low, with the exception of cases with large fleets. This is because car trips are more sensitive to waiting times and larger fleets have generally lower waiting times (see Figure 13). This sensitivity to waiting time is due to cars being priced similarly to SAV (in the full cost scenario) and being generally readily available to the driver. Predictably, the ‘full cost’ scenario has a much higher shift from cars. This in turn lowers the shift from all other modes since more of the vehicles are

³Pricing of one-way free floating car-sharing systems are generally per minute, and 0.5 €/km is equivalent to 0.25 €/min when driving at an average speed of 30 km/h.

Table 2. Results for different fleet sizes and prices per km with predictive relocation

Case	Fleet size	Price €/km	Modal shift (%)					Fleet results			
			Total	Bicycle	Car	Foot interzonal	PT	VT	VKT	RV €	eVKT %
Marginal cost	2500	0.5	4.38	0.71	0.18	12.06	8.01	74.2	704	334	14.84
		1.0	5.23	0.50	0.29	16.28	8.15	77.4	596	524	15.66
		2.0	6.32	0.87	0.78	20.38	8.26	89.6	433	733	17.26
	5000	0.5	8.91	1.86	0.50	19.01	20.47	74.8	599	319	11.54
		1.0	10.84	1.44	0.91	28.81	20.29	79.8	491	478	12.17
		2.0	9.47	1.59	1.44	30.65	11.58	67.0	301	527	14.80
	10000	0.5	17.28	5.17	1.60	29.86	43.27	71.3	499	296	9.19
		1.0	20.00	4.60	3.21	47.76	37.77	71.8	374	411	9.73
		2.0	10.96	2.00	1.81	35.48	13.04	40.3	176	302	14.21
	20000	0.5	30.58	14.95	6.12	48.40	69.66	56.2	361	249	7.68
		1.0	26.24	8.50	6.16	59.90	46.44	44.0	218	261	8.93
		2.0	11.49	2.15	1.96	37.18	13.53	21.7	95	158	15.83
Full cost	2500	0.5	3.67	0.32	3.01	8.93	3.33	70.2	868	407	12.08
		1.0	5.37	0.39	1.68	15.46	7.29	80.7	630	544	15.78
		2.0	6.54	0.80	2.21	19.37	7.88	92.4	438	745	17.08
	5000	0.5	7.23	0.73	7.07	13.24	8.40	66.7	761	381	9.43
		1.0	11.19	1.11	4.97	26.38	17.63	83.6	526	503	11.96
		2.0	10.25	1.52	4.10	29.97	11.38	69.4	306	557	14.30
	10000	0.5	14.22	1.82	15.67	19.00	19.34	62.0	653	345	7.78
		1.0	21.16	3.06	13.91	41.37	32.46	79.3	427	450	9.09
		2.0	12.13	1.98	5.19	35.26	12.94	42.2	181	325	13.79
	20000	0.5	27.54	4.77	32.79	28.38	39.65	57.4	539	301	6.97
		1.0	33.14	7.15	28.87	56.56	44.20	54.5	281	343	7.76
		2.0	12.83	2.16	5.61	37.17	13.52	22.9	99	172	14.60

*VT= vehicle-trips; VKT= vehicle-km traveled; RV= revenue per vehicle (€); eVKT= empty VKT share (%). All quantities are per day averages.

used to serve former drivers. Bicycle trips have some of the lowest rate of modal shift, demonstrating that bicycles are a very effective way to move in the city. This of course may be influenced by weather and season, which was not considered in the survey nor in our model. In other words, it is conceivable that the modal shift of cycling to SAV will fluctuate significantly over a year. A more alarming result is that related to potential behavioral changes of current public transport users. Modal shifts from PT trips are some of the highest, especially in the ‘marginal cost’ case, touching nearly 70% for the case with a fleet of 20,000 vehicles and a price of 0.5 €/km. Even in the full cost case and with a more realistic fleet of 5,000 vehicles, shift from PT is between 8% and 17%. This would severely impact PT revenue and service, thus possibly triggering a negative spiral. Less demand could lead to service reductions and make it even less attractive to use public transport. However, in order to study the longer term effects on PT, a different modeling approach is needed. As noted, our binary choice model of current mode versus SAV is based on the assumption that the attributes of the original modes do not change.

Figure 12a shows the total modal shift for the full car cost case with predictive relocation as a function of fleet size. As expected, modal shift tends to increase with fleet size. However, with more expensive prices a plateau is reached, in our case with a fleet of about 10,000 vehicles. Figure 12b reports the same results but as a function of prices. Interestingly, modal shift does not necessarily increase with lower prices. For




Figure 10. Potential demand of SAV by original mode (log scale). Case with 1 €/km fare and predictive relocation. Results for marginal and full car costs are identified with dashed and solid lines, respectively.




Figure 11. (a) Trips per vehicle per day and vehicle-km traveled per day; (b) revenue-vehicle per day (€) and empty vehicle-km traveled (%). Case with 1 €/km fare and predictive relocation. Dashed lines and x represent the marginal cost case, while solid lines and circles represent the full cost case.

larger fleets, there is a price that maximizes modal shift, and a lower price actually decreases the shift. This is because lower prices per km tend to attract longer trips that keep vehicles occupied for longer, thus decreasing their ability to serve more passengers. Note that vehicle-km traveled (VKT) instead always increase with lower prices.

The average waiting time is generally relatively low, especially for larger fleets. Figure 13 shows average waiting times for passengers shifting from each mode. The relatively low average waiting time is a result of passengers having an alternative to the system in their original mode.

Figure 14 shows the origins of trips by mode and the relative modal shift for each mode by origin location. Modal shift from bicycle happens mostly at peripheral nodes. Car shift is concentrated in the center. Shift from foot (walking trips) and public transport (PT) is more homogeneous, although higher in the periphery.

Figure 15 highlights the modal shift trends when grouping origin and destination by relative popularity. The trends are quite different depending on the original mode considered. The largest shift from bicycles happen when trips are between the least popular nodes. This is line with the general trend for total shift. Car drivers are more likely to shift to SAV when traveling between nodes of similar popularity. Walking trips have the opposite trend. These trends tend to reflect the relative distance of trips that shift from each mode: longer walking trips tend to have higher modal shift




Figure 12. Total modal shift as a function of: a) fleet size for different prices; b) price for different fleet sizes.




Figure 13. Average waiting time for trips shifting from each mode. Case with 1 €/km fare and predictive relocation.

and be between areas of widely different popularity (e.g. from periphery to the center). The opposite is true for car trips. As expected, shifts from PT trips tend to increase as the popularity of origin and destination node decreases.

Table 3 reports the results for the same scenarios in the case without predictive relocation. In this case, vehicles just move only in response to a specific request, not in advance. Due to the absence of predictive actions, waiting time in some cases can be higher because the closest vehicle can be far away when a request arrives. Without predictive relocation the modal shift is lower, with correspondingly lower revenue per vehicle. However, the system has less empty vehicle-km traveled. Predictive relocation is especially effective when prices are high: in the case with the highest price of €2/km, the total mode shift increases by between 13.9% and 28.8% with predictive relocation. Similarly, revenue per vehicle also increase by between 12.2% and 26.4% for the cases with the highest price, suggesting that relocation is an effective strategy to increase revenues for the system operator. For cases with low prices and small fleets, relocation has almost no effect. This is because the system is not constrained by demand asymmetry, but by availability of vehicles. Introducing relocation therefore does not significantly increase the amount of demand that can be served by the system. The mode most affected by introducing relocation is the private car: mode shift from car increases by as much as 46.9% (for the case with €2/km, fleet size of 5000 vehicles and marginal cost).




Figure 14. Maps of origins by mode (top) and modal shift by origin (bottom). Both scales are as a fraction of the 99th percentile value for each case (to exclude outliers). Case with full car cost, 5000 SAVs, 1 €/km fare, and predictive relocation.




Figure 15. Modal shift by mode and origin/destination grouped in quintiles by popularity (number of origins and destinations in each node). Case with full car cost, predictive relocation, fleet size of 5000 vehicles and 1 €/km fare. The first and second quintile are merged due to the low number of trips.

6.2. Sensitivity analysis

Here we report the results of our sensitivity analysis on our model. For these simulations we use a baseline scenario (5000 vehicles, 1 €/km, and full cost assumption) and we run our model only on one day (Monday) to have lower running times.

6.2.1. Minimum price

The effect of the minimum price on modal shift is shown in Figure 16a. Car and bike trips are most sensitive to the minimum price, with the shift decreasing by almost 60% when the minimum price increases from 0 to 5 €/km. Walking trips decrease by 34%. Public transport trips are the least affected, without significant change. This is because PT trips tend to be longer (see Fig. 8), thus are the most likely to already induce a significant fare and are hence least affected by a raise in minimum price. Therefore, PT trips are the most likely to take advantage of the unused system capacity resulting from the decreased shift from other modes.

6.2.2. Value of time

Figure 16b shows the effect of value of time (VOT) on the modal shift. Walking trips are more likely to shift to SAV when we assume a higher VOT. This is because the walking trips have the lowest average speed, and thus a shift to SAV can save the most time. Moreover, since walking has no monetary costs, the time cost component is large. Conversely, shift from car trips (which are the opposite of walking trips in both respects) decreases substantially with increasing value of time. Modal shift from public transport and bike trips show no strong correlation with value of time. An

Table 3. Results for different fleet sizes and prices per km without predictive relocation

Case	Fleet size	Price €/km	Modal shift (%)					Fleet results			
			Total	Bicycle	Car	Foot interzonal	PT	VT	VKT	RV €	eVKT %
Marginal cost, no relocation	2500	0.5	4.34	0.67	0.16	12.02	7.94	73.6	677	330	13.17
		1.0	5.19	0.47	0.28	16.22	8.06	76.7	575	520	13.79
		2.0	5.50	0.71	0.64	17.78	7.29	81.4	390	650	16.06
	5000	0.5	8.82	1.77	0.46	18.91	20.35	73.7	577	315	9.88
		1.0	10.71	1.39	0.88	28.57	20.01	78.9	477	474	10.27
		2.0	7.35	1.09	0.98	23.69	9.44	48.4	223	417	13.42
	10000	0.5	17.07	4.95	1.50	29.59	42.97	69.9	479	292	7.32
		1.0	18.16	3.68	2.56	44.13	34.68	67.5	348	378	7.70
		2.0	8.83	1.43	1.31	28.49	11.00	28.5	126	245	11.32
	20000	0.5	28.89	13.31	5.23	45.98	67.28	53.3	335	236	5.32
		1.0	22.54	6.02	4.28	52.69	41.69	37.4	186	227	6.19
		2.0	9.79	1.68	1.53	31.58	11.97	16.2	70	136	9.79
Full cost, no relocation	2500	0.5	3.61	0.27	2.95	8.88	3.24	69.0	841	401	10.82
		1.0	5.33	0.39	1.65	15.43	7.24	80.3	601	539	13.48
		2.0	5.74	0.66	1.86	17.06	7.02	82.8	390	664	15.91
	5000	0.5	7.14	0.65	6.97	13.17	8.29	65.6	740	376	8.26
		1.0	11.09	1.07	4.87	26.27	17.49	82.7	504	499	9.76
		2.0	7.96	1.07	2.90	23.38	9.31	51.4	231	441	13.01
	10000	0.5	14.03	1.65	15.47	18.83	19.14	60.8	635	341	6.50
		1.0	20.12	2.79	12.78	39.99	31.13	76.5	402	430	6.82
		2.0	9.70	1.42	3.83	28.29	10.92	29.5	128	263	10.89
	20000	0.5	27.08	4.39	32.28	27.96	39.17	55.9	520	297	5.32
		1.0	28.29	5.24	22.47	50.69	40.04	48.7	245	295	5.07
		2.0	10.82	1.67	4.45	31.47	11.91	17.2	72	145	9.37

*VT= vehicle-trips; VKT= vehicle-km traveled; RV= revenue per vehicle (€); eVKT= empty VKT share (%). All quantities are per day averages.

interesting result is that the value of time does not influence the total modal shift of the system, but just the relative share of the shift by original mode.

6.2.3. Cost of car

We also studied how the car cost affects the results (Figure 16c). We considered additional four cost scenarios centered around the assumed cost of 0.72 €/km, at $\pm 20\%$ and $\pm 50\%$. The cost of car has a significant effect on the results. For costs higher than about 1 €/km, the shift from cars is actually higher than from public transport trips.

6.2.4. Cost of public transport

Figure 16d shows the sensitivity of PT modal shift to monetary cost and travel time, as a percentage of the base scenario. Travel time in particular has a strong influence on the results: a decrease of 20% from the base scenario gives a 35% decrease in modal shift from PT. For the same relative change in fares, the decrease is only 16%. This suggests that changes in travel time have a much stronger influence on mode choice.




Figure 16. Effect on modal shift of a) minimum tariff; b) value of time (VOT); c) car cost per km; d) PT cost and travel time. Results for a fleet of 5,000 SAV with predictive relocation and a price of 1 €/km.

6.3. Discussion and policy implications

Our results lead to several policy implications. We show how SAVs have the potential to capture a large share of mobility, both at the urban and suburban level. This could have enormous benefits, such as reducing private car ownership, reducing car parking requirements, and increasing accessibility. However, our results also suggest that SAVs are likely leading to a major shift from current public transport users, making this mode economically unsustainable and possibly triggering a deterioration of service which would lead to even bigger modal shifts.

We furthermore show the potential of shifts from active modes to SAVs. Quantifying this shift is arguably associated with more uncertainty as health benefits and other aspects that are difficult to capture with utility play a larger role for choosing cycling and walking. Nevertheless, our results suggest that this possibly undesired shift can not be ignored. Both modal shifts will increase the road traffic and could lead to higher traffic congestion which transport planners need to prepare for. Our economic analysis show that SAVs are profitable even at small fleet sizes and high prices, suggesting that deployment could be quick once the technology becomes widely available. Moreover, larger fleets almost always lead to much better economic outcomes for the fleet operator, and the biggest profits coincide with the largest mode shifts. Unregulated SAVs could therefore quickly take over a large portion of the transportation in a city, and thus could lead to major negative effects. It is therefore essential to introduce regula-

tion that internalize costs related to congestion, and allow SAVs to work in synergy with other modes. A possible solution for the first point is introducing a road tax or other form of congestion tax to account for the physical space that these vehicles take in the city. For the second point, a better integration with the public transport system may lead to higher efficiency for both modes, with public transport focusing on major corridors while leaving the more expensive connections in low-density areas to SAVs. Enabling of multi-modal trips using SAVs could therefore be beneficial towards a better mobility system.

7. Conclusions

Shared autonomous vehicles have the potential to deeply change future transport systems. In this work we introduce a framework to investigate the potential modal shift induced by this mode through a simulation model and a trip generation model trained on transport survey data. We present results for a case study with the metropolitan area of Munich, Germany. Our trip generation model is based on a Bayesian approach of OD demand generation that allows us to run a large number of scenarios simulating the city at a full scale over several days with realistic OD flow patterns. Our scalable SAV simulation model allows to study the potential SAV demand from different modes and different trip patterns considering SAV fleet properties, by taking into account the resulting real-time waiting times and vehicle locations, which influence traveler’s mode choice.

Our results illustrate how the SAV demand is highly dependent on the fleet size and price of the SAV system. We also show that the cost perception of the private car is an important parameter to estimate the potential demand. Generally, it is desired by policy makers that users appreciate the full cost of their car usage. We show that it might also impact the resulting appreciation of an SAV alternative. Whether the ‘full cost’ scenarios in our simulation would, in the long run, indeed lead to users abandoning their private car is though beyond the scope of our study as we only study trips and do not consider tours and family interactions. We do show, however, that more SAV usage by private car users due to higher cost perception can also reduce the SAV demand from public transport users if the SAV fleet is limited. This might be another welcome effect of an increase in (perceived) costs for private car usage.

The analysis of economic performance suggests that operators are likely to favor larger fleets and prices higher than current car sharing but much lower than taxis. The results from this work demonstrate the potential for SAVs to capture a large portion of urban and suburban mobility, and the economic viability of very large SAV fleets. Unregulated SAVs could reduce private car usage, but also make public transport unsustainable due to the decreases in passenger levels, especially outside of city centers. Although not considered in this work, large fleets of SAVs are also likely to significantly increase congestion. We close by reminding about limitations and pointing out further work directions, some of which have already been mentioned throughout the manuscript. Congestion modelling and tour instead of trip based modelling would likely lead to an SAV demand that is lower than the ones reported in our study. One could modify our approach to include these aspects at the likely cost of reducing the ability to model the whole metropolitan area in an efficient way as carried out in this research. We therefore suggest rather a direction of further work could be to compare existing agent-based models and our proposed approach for smaller networks. Overall, we envision that the framework proposed here could guide modellers to then study

specific parts of the network or population groups in more detail for which a large SAV shift might be expected and the above noted limitations are significant. We conclude by noting a yet different promising work direction: We did not consider electrified vehicles in this work. Considering charging constraints as well as possibilities for grid-services provided by SAEVs is one of our next work directions.

References

- Aboulela, M., Al Haddad, C., & Antoniou, C. (2021, June). Are young users willing to shift from carsharing to scooter-sharing? *Transportation Research Part D: Transport and Environment*, 95, 102821. Retrieved 2021-10-11, from <https://linkinghub.elsevier.com/retrieve/pii/S1361920921001243>
- Agora Verkehrswende. (2019). *Parkraummanagement lohnt sich! Leitfaden für Kommunikation und Verwaltungspraxis*.
- Alam, M. J., & Habib, M. A. (2018). Investigation of the Impacts of Shared Autonomous Vehicle Operation in Halifax, Canada Using a Dynamic Traffic Microsimulation Model. *Procedia Computer Science*, 130, 496–503. Retrieved 2021-09-29, from <https://linkinghub.elsevier.com/retrieve/pii/S1877050918304265>
- Amanatidis, T., Langdon, P., & Clarkson, P. J. (2018, March). Needs and Expectations for Fully Autonomous Vehicle Interfaces. In *Companion of the 2018 ACM/IEEE International Conference on Human-Robot Interaction* (pp. 51–52). Chicago IL USA: ACM. Retrieved 2021-09-29, from <https://dl.acm.org/doi/10.1145/3173386.3177054>
- Axsen, J., & Sovacool, B. K. (2019, June). The roles of users in electric, shared and automated mobility transitions. *Transportation Research Part D: Transport and Environment*, 71, 1–21. Retrieved 2021-10-13, from <https://www.sciencedirect.com/science/article/pii/S1361920918312483>
- Bauer, G. S., Greenblatt, J. B., & Gerke, B. F. (2018, April). Cost, Energy, and Environmental Impact of Automated Electric Taxi Fleets in Manhattan. *Environmental Science & Technology*, 52(8), 4920–4928.
- Bösch, P. M., Becker, F., Becker, H., & Axhausen, K. W. (2018, May). Cost-based analysis of autonomous mobility services. *Transport Policy*, 64, 76–91. Retrieved 2021-10-13, from <https://www.sciencedirect.com/science/article/pii/S0967070X17300811>
- Cao, Z., Zhang, X., Chua, K., Yu, H., & Zhao, J. (2021, May). E-scooter sharing to serve short-distance transit trips: A Singapore case. *Transportation Research Part A: Policy and Practice*, 147, 177–196. Retrieved 2021-10-11, from <https://linkinghub.elsevier.com/retrieve/pii/S0965856421000598>
- Carron, A., Seccamonte, F., Ruch, C., Frazzoli, E., & Zeilinger, M. N. (2021, March). Scalable Model Predictive Control for Autonomous Mobility-on-Demand Systems. *IEEE Transactions on Control Systems Technology*, 29(2), 635–644. Retrieved 2021-09-28, from <https://ieeexplore.ieee.org/document/8938759/>
- Correia, G. H. d. A., & van Arem, B. (2016, May). Solving the User Optimum Privately Owned Automated Vehicles Assignment Problem (UO-POAVAP): A model to explore the impacts of self-driving vehicles on urban mobility. *Transportation Research Part B: Methodological*, 87, 64–88.
- Eom, J.-K., Lee, K.-S., Ko, S., & Lee, J. (2022, January). Exploring Travel Mode Preference of External Trips for Smart City Transportation Planning: Sejong, Korea. *Sustainability*, 14(2), 630. Retrieved 2022-12-23, from <https://www.mdpi.com/2071-1050/14/2/630>
- Estandia, A., Schiffer, M., Rossi, F., Luke, J., Kara, E. C., Rajagopal, R., & Pavone, M. (2021, February). On the Interaction between Autonomous Mobility on Demand Systems and Power Distribution Networks – An Optimal Power Flow Approach. *arXiv:1905.00200 [cs, eess]*. Retrieved 2021-02-19, from <http://arxiv.org/abs/1905.00200> (arXiv: 1905.00200)

- Etminani-Ghasrodashti, R., Ketankumar Patel, R., Kermanshachi, S., Rosenberger, J. M., & Foss, A. (2022, May). Modeling Users' Adoption of Shared Autonomous Vehicles Employing Actual Ridership Experiences. *Transportation Research Record: Journal of the Transportation Research Board*, 036119812210936. Retrieved 2022-06-07, from <http://journals.sagepub.com/doi/10.1177/0361198122109362>
- Fagnant, D. J., & Kockelman, K. M. (2014, March). The travel and environmental implications of shared autonomous vehicles, using agent-based model scenarios. *Transportation Research Part C: Emerging Technologies*, 40, 1–13.
- Follmer, R., & Gruschwitz, D. (2019). *Mobility in Germany - short report*. Bundesministerium für verkehr und digitale infrastruktur. Retrieved from http://www.mobilitaet-in-deutschland.de/pdf/MiD2017_ShortReport.pdf
- Gao, K., Shao, M., Axhausen, K. W., Sun, L., Tu, H., & Wang, Y. (2021, July). Inertia effects of past behavior in commuting modal shift behavior: interactions, variations and implications for demand estimation. *Transportation*. Retrieved 2021-10-11, from <https://link.springer.com/10.1007/s11116-021-10203-6>
- Gelauff, G., Ossokina, I., & Teulings, C. (2019, March). Spatial and welfare effects of automated driving: Will cities grow, decline or both? *Transportation Research Part A: Policy and Practice*, 121, 277–294. Retrieved 2021-10-13, from <https://www.sciencedirect.com/science/article/pii/S0965856417311758>
- Golbabaei, F., Yigitcanlar, T., & Bunker, J. (2021, August). The role of shared autonomous vehicle systems in delivering smart urban mobility: A systematic review of the literature. *International Journal of Sustainable Transportation*, 15(10), 731–748. Retrieved 2021-09-25, from <https://www.tandfonline.com/doi/full/10.1080/15568318.2020.1798571>
- Haboucha, C. J., Ishaq, R., & Shiftan, Y. (2017, May). User preferences regarding autonomous vehicles. *Transportation Research Part C: Emerging Technologies*, 78, 37–49. Retrieved 2022-06-07, from <https://linkinghub.elsevier.com/retrieve/pii/S0968090X17300177>
- Heilig, M., Hilgert, T., Mallig, N., Kagerbauer, M., & Vortisch, P. (2017). Potentials of Autonomous Vehicles in a Changing Private Transportation System – a Case Study in the Stuttgart Region. *Transportation Research Procedia*, 26, 13–21. Retrieved 2021-11-09, from <https://linkinghub.elsevier.com/retrieve/pii/S2352146517308633>
- Huang, Y., Kockelman, K. M., Garikapati, V., Zhu, L., & Young, S. (2021, February). Use of Shared Automated Vehicles for First-Mile Last-Mile Service: Micro-Simulation of Rail-Transit Connections in Austin, Texas. *Transportation Research Record: Journal of the Transportation Research Board*, 2675(2), 135–149. Retrieved 2022-12-23, from <http://journals.sagepub.com/doi/10.1177/0361198120962491>
- Hörl, S., Becker, F., & Axhausen, K. W. (2021, February). Simulation of price, customer behaviour and system impact for a cost-covering automated taxi system in Zurich. *Transportation Research Part C: Emerging Technologies*, 123, 102974. Retrieved 2022-12-20, from <https://linkinghub.elsevier.com/retrieve/pii/S0968090X21000115>
- Hörl, S., Ruch, C., Becker, F., Frazzoli, E., & Axhausen, K. W. (2019, May). Fleet operational policies for automated mobility: A simulation assessment for Zurich. *Transportation Research Part C: Emerging Technologies*, 102, 20–31. Retrieved 2021-10-11, from <https://www.sciencedirect.com/science/article/pii/S0968090X18304248>
- Iacobucci, R., Bruno, R., & Schmöcker, J.-D. (2021, June). An Integrated Optimisation-Simulation Framework for Scalable Smart Charging and Relocation of Shared Autonomous Electric Vehicles. *Energies*, 14(12), 3633. Retrieved 2021-06-22, from <https://www.mdpi.com/1996-1073/14/12/3633>
- Iacobucci, R., Donhauser, J., Schmöcker, J., & Pruckner, M. (2021). Frequency Control Reserve Provision from a Fleet of Shared Autonomous Electric Vehicles. In *7th IEEE International Conference on Models and Technologies for Intelligent Transportation Systems (MT-ITS 2021)*.
- Iacobucci, R., McLellan, B., & Tezuka, T. (2018, June). Modeling shared autonomous electric vehicles: Potential for transport and power grid integration. *Energy*.

- Iacobucci, R., McLellan, B., & Tezuka, T. (2019, March). Optimization of shared autonomous electric vehicles operations with charge scheduling and vehicle-to-grid. *Transportation Research Part C: Emerging Technologies*, 100, 34–52. Retrieved 2020-05-19, from <https://linkinghub.elsevier.com/retrieve/pii/S0968090X18309197>
- Iacobucci, R., Pruckner, M., & Schmöcker, J.-D. (2022, September). A large scale simulation of the electrification effects of SAVs. Skiathos, Greece.
- INRIX. (2017). *The Impact of Parking Pain in the US, UK and Germany*.
- LaMondia, J. J., Fagnant, D. J., Qu, H., Barrett, J., & Kockelman, K. (2016). Long-distance travel mode shifts due to automated vehicles: a statewide mode-shift simulation experiment and travel survey analysis. In *Presented at the 95th Annual Meeting of the Transportation Research Board* (Vol. 1, p. 3). (Issue: 2)
- Letmathe, P., & Soares, M. (2017, December). A consumer-oriented total cost of ownership model for different vehicle types in Germany. *Transportation Research Part D: Transport and Environment*, 57, 314–335. Retrieved 2021-09-01, from <https://linkinghub.elsevier.com/retrieve/pii/S1361920917304868>
- Liao, Z., Taiebat, M., & Xu, M. (2021, November). Shared autonomous electric vehicle fleets with vehicle-to-grid capability: Economic viability and environmental co-benefits. *Applied Energy*, 302, 117500. Retrieved 2021-09-25, from <https://linkinghub.elsevier.com/retrieve/pii/S0306261921008850>
- Loeb, B., & Kockelman, K. M. (2019, March). Fleet performance and cost evaluation of a shared autonomous electric vehicle (SAEV) fleet: A case study for Austin, Texas. *Transportation Research Part A: Policy and Practice*, 121, 374–385. Retrieved 2021-10-11, from <https://www.sciencedirect.com/science/article/pii/S096585641730112X>
- Loeb, B., Kockelman, K. M., & Liu, J. (2018). Shared autonomous electric vehicle (SAEV) operations across the Austin, Texas network with charging infrastructure decisions. *Transportation Research Part C: Emerging Technologies*, 89.
- Matalqah, I., Shatanawi, M., Alatawneh, A., & Mészáros, F. (2022, December). Impact of Different Penetration Rates of Shared Autonomous Vehicles on Traffic: Case Study of Budapest. *Transportation Research Record: Journal of the Transportation Research Board*, 2676(12), 396–408. Retrieved 2022-12-23, from <http://journals.sagepub.com/doi/10.1177/03611981221095526>
- Melendez, K. A., Das, T. K., & Kwon, C. (2020, December). Optimal operation of a system of charging hubs and a fleet of shared autonomous electric vehicles. *Applied Energy*, 279, 115861. Retrieved 2020-12-02, from <https://linkinghub.elsevier.com/retrieve/pii/S0306261920313349>
- Mo, B., Cao, Z., Zhang, H., Shen, Y., & Zhao, J. (2021, June). Competition between shared autonomous vehicles and public transit: A case study in Singapore. *Transportation Research Part C: Emerging Technologies*, 127, 103058. Retrieved 2021-09-25, from <https://linkinghub.elsevier.com/retrieve/pii/S0968090X21000863>
- Mo, B., Wang, Q. Y., Moody, J., Shen, Y., & Zhao, J. (2021, September). Impacts of subjective evaluations and inertia from existing travel modes on adoption of autonomous mobility-on-demand. *Transportation Research Part C: Emerging Technologies*, 130, 103281. Retrieved 2021-09-25, from <https://linkinghub.elsevier.com/retrieve/pii/S0968090X2100293X>
- Mobileye Inc. (2021, September). *Launching our Self-Driving Robotaxi in Germany*. Retrieved 2021-09-07, from <https://www.mobileye.com/blog/self-driving-robotaxi-sixt-germany-iaa/>
- Münchner Verkehrs- und Tarifverbund (MVG). (n.d.). *MVV Tickets & Day Tickets*. Retrieved 2021-07-14, from <https://www.mvv-muenchen.de/en/tickets-and-fares>
- Münchner Verkehrsgesellschaft mbH (MVG). (n.d.). *MVG in figures - Key facts*. Münchner Verkehrsgesellschaft mbH (MVG).
- Narayanan, S., Chaniotakis, E., & Antoniou, C. (2020, February). Shared autonomous vehicle services: A comprehensive review. *Transportation Research Part C: Emerging Technologies*, 111, 255–293.

- Pruckner, M., & Eckhoff, D. (2020, October). Shared Autonomous Electric Vehicles and the Power Grid: Applications and Research Challenges. In *2020 IEEE PES Innovative Smart Grid Technologies Europe (ISGT-Europe)* (pp. 1151–1155).
- Schubert, E., & Rousseeuw, P. J. (2019). Faster k-Medoids Clustering: Improving the PAM, CLARA, and CLARANS Algorithms. In *Similarity Search and Applications* (Vol. 11807, pp. 171–187). Cham: Springer International Publishing. Retrieved 2021-09-28, from http://link.springer.com/10.1007/978-3-030-32047-8_16 (Series Title: Lecture Notes in Computer Science)
- Shen, Y., Zhang, H., & Zhao, J. (2018, July). Integrating shared autonomous vehicle in public transportation system: A supply-side simulation of the first-mile service in Singapore. *Transportation Research Part A: Policy and Practice*, 113, 125–136. Retrieved 2021-09-29, from <https://linkinghub.elsevier.com/retrieve/pii/S096585641730681X>
- Silver, D. (2021, January). AutoX Launches Driverless Robotaxis To The Public In Shenzhen. *Forbes*. Retrieved 2021-02-12, from <https://www.forbes.com/sites/davidsilver/2021/01/27/autox-launches-driverless-robotaxis-in-shenzhen/>
- Sprei, F. (2018, March). Disrupting mobility. *Energy Research & Social Science*, 37, 238–242. Retrieved 2021-10-13, from <https://www.sciencedirect.com/science/article/pii/S2214629617303596>
- Stathopoulos, A., Cirillo, C., Cherchi, E., Ben-Elia, E., Li, Y.-T., & Schmöcker, J.-D. (2017, December). Innovation adoption modeling in transportation: New models and data. *Journal of Choice Modelling*, 25, 61–68. Retrieved 2021-10-13, from <https://www.sciencedirect.com/science/article/pii/S1755534517300210>
- Stoiber, T., Schubert, I., Hoerler, R., & Burger, P. (2019, June). Will consumers prefer shared and pooled-use autonomous vehicles? A stated choice experiment with Swiss households. *Transportation Research Part D: Transport and Environment*, 71, 265–282. Retrieved 2022-12-23, from <https://linkinghub.elsevier.com/retrieve/pii/S1361920918303304>
- Tian, Z., Feng, T., Timmermans, H. J., & Yao, B. (2021, July). Using autonomous vehicles or shared cars? Results of a stated choice experiment. *Transportation Research Part C: Emerging Technologies*, 128, 103117. Retrieved 2022-12-23, from <https://linkinghub.elsevier.com/retrieve/pii/S0968090X21001364>
- Vosooghi, R., Puchinger, J., Jankovic, M., & Vouillon, A. (2019, October). Shared autonomous vehicle simulation and service design. *Transportation Research Part C: Emerging Technologies*, 107, 15–33. Retrieved 2022-12-23, from <https://linkinghub.elsevier.com/retrieve/pii/S0968090X19304449>
- White, J. (2020, October). Waymo opens driverless robo-taxi service to the public in Phoenix. *Reuters*. Retrieved 2021-02-12, from <https://www.reuters.com/article/us-waymo-autonomous-phoenix-idUKKBN26T2Y3>
- Zhang, R., & Pavone, M. (2016, January). Control of robotic mobility-on-demand systems: A queueing-theoretical perspective. *The International Journal of Robotics Research*, 35(1-3), 186–203. Retrieved 2019-06-06, from <https://doi.org/10.1177/0278364915581863>
- Zhang, R., Rossi, F., & Pavone, M. (2016, May). Model predictive control of autonomous mobility-on-demand systems. In *2016 IEEE International Conference on Robotics and Automation (ICRA)* (pp. 1382–1389).

S-I. Supplementary material

S-I.1. Incomplete data pre-augmentation

The spatial information about the trips (origin and destination) in the MiD dataset is given in a grid structure with square grid cells, available in three different resolutions: 500m, 1km, 5km. However, not every resolution is available for every trip. About 50% of trips lack 1 km information, and roughly 25% has incomplete 5 km information. As the data availability is getting lower with finer resolution, we decided to use the 1km representation for this work, where origin and destination is given in 341,408 cases.

We also need to take into consideration those grid cells which have no trips within the dataset due to low trip frequency. For this reason, adding some diffusion to the model might be desirable. The incomplete data should not be ignored because it does not affect all modalities equally, and thus its omission may bias the resulting model. We found that this problem occurs primarily (but not exclusively) when dealing with short trips. Since there is a strong dependency between the chosen modality and the distance of a trip (see Figure S6), ignoring trips with incomplete spatial information would cause a significant shift in the modality patterns, i.e., less trips by foot and bicycle. To resolve the modality shift problem we artificially create suitable cells depending on the known data. In this section we discuss an approach to estimate missing grid cell data while preserving the mode patterns. Most of the trips with a missing coordinate (either origin or destination) nonetheless have an associated trip distance. This allows us to guess the missing coordinate within a given radius of the coordinate present.


To make the most data reusable, a augmentation consisting of multiple steps is performed, depending on the availability of living place, 5 km cells, trip length.

The following operations are performed:

- If no information is given, i.e. no living place, no 5km origin- or destination-cell and no trip-length is given, the trip can not be used. This applies only to a relatively small fraction of the data (less than 10%).
- If only the living place is given, we sample trips around the living place using the approximate trip-length. This is done by sampling a distance from a Rayleigh distribution⁴. The artificial cell is now chosen from all cells with the correct distance to the known cell.
- If only one 5km cell is given, we estimate the other cell using the known trip-length like above
- If complete 5km information is given:
 - If no 1km information is given, we select a random 1km cell within the 5km cell.
 - Estimate the second 1km-cell using the known trip-length as described above.

This approach tends to increase the diffusion due to the randomness introduced. Because of this effect, the augmented data is excluded for the spatial KDE in Section S-I.2.2.

⁴The Rayleigh distribution gives the distance between two normal distribution points in 2d. The Rayleigh distribution is a special case of the Chi distribution.



(a) Spatial distribution of training-data with known cells (b) Spatial distribution of pre-augmented training data

Figure S1. Comparison of the spatial distribution of data with and without the pre-augmentation. Both graphs are in logarithmic scale. Since the pre-augmentation adds diffusion, the affected data is excluded from the KDE fitting in Section S-I.2.2

S-I.2. Continuous space/time domain

In this section we explain the leveraging from discrete clusters to continuous coordinates and the fitting of a continuous time using Gaussian⁵ Kernel Density Estimation (KDE) and Gaussian Mixture Models (GMMs).

Both KDE and GMM methods utilize Gaussian Kernels to fit a distribution to known data (N samples), but in a different way: KDE is taking a sum of N equally weighted Gaussians with equal window size ε , where each sample is wrapped by a Gaussian kernel; GMMs try to fit a distribution (given by the samples) using $n_{kernels}$ individually weighted Gaussians with adaptable window sizes. KDEs depend on the windows-size hyper-parameter ε , while $n_{kernels}$ is a hyper-parameter for GMMs. Further GMMs are only applicable under the assumption that the searched distribution is a combination normal distributions. Figure S3 shows the fitted GMMs for each week day and each purpose in our model.

Formula for a GMM with $n_{kernels}$ kernels:

$$\sum_{k=1}^{n_{kernels}} \phi_k \mathcal{N}(\mu_k, \Sigma_k) \quad | \quad \phi^T \mathbf{1} = 1 \quad (S1)$$

μ_k is the mean of all samples belonging to cluster k , Σ_k the respective covariance matrix.

⁵There are other Kernels possible, but for the sake of simplicity we restrict ourselves to Gaussians.




Figure S2. a) weighted combination of all cluster-KDEs which have been fitted to the training data; b) a histogram of 20,000 sampled agents. Both figures are in log scale. The model is quite good at capturing the shape of the city, but it still has some diffusion. Forested areas around the city, which are shown in gray, are unlikely to spawn agents in general. The overall shape is similar to the distribution of known data in Figure S1a.

Formula for KDE:

$$\frac{1}{N} \sum_{i=1}^N \mathcal{N}(x_i, \epsilon) \quad (\text{S2})$$

where $x_1, \dots, x_N$ denotes the samples.

S-I.2.1. Time fitting

As can be seen in Figure 4, the time in our model is conditioned by the weekday, the origin cluster, and the purpose, which are all categorical values. Thus, for each possible combination of those three attributes, the data will be handled separately. A direct consequence of this is that the size of the data chunks is relatively small and thus prone to overfitting.

While using KDE for the time fitting would be possible from a technical perspective, in general KDE is not very robust to overfitting since each sample is associated with one kernel. What we found is that, in contrast to KDE, a GMM with few degrees of freedom can better deal with limited data.

In our case we chose $n_{\text{kernel}s} = 3$, which we found to be sufficient to capture the time distribution under a given trip-purpose. A smaller number of kernels did not suffice, while a bigger number did not give any improvements. The model could be improved in the future by implementing different $n_{\text{kernel}s}$ for different trip-purposes.

It is important to add that the day-time t (in hours) is cyclic, i.e. the range of possible values is $T := [0, 24)$ and not $\mathbb{R}$. So in theory instead of fitting a regular GMM (on $\mathbb{R}$), one would have to fit a wrapped GMM, which is limited to the range T . In our case this issue is not relevant from a practical perspective, as the amount of traffic during the night hours is negligible. This low amount of data at the boundary

of the time range T makes both fittings similar.

S-I.2.2. Space fitting

In an idealized scenario with a very large amount of data available, the BN would directly sample continuous origin-destination pairs from a 4 dimensional distribution, with 2 dimensions each for both origin and destination. The BN in this work samples origin and destination clusters instead. So instead of using one global 4 dimensional origin-destination estimator, we build local estimators for the clusters. Depending on data availability, we use one of two different approaches for clusters A and B:

- If sufficient data is seen from A to B: In general, this case is likely to happen in dense neighboring clusters. Since the data availability is good, a KDE may be fitted to the respective origin-destination pairs.
- If the number of trips from A to B is low: The number of training points is not appropriate for a 4 dimensional space. In this case the dependency between the exact origin and destination is dropped⁶ and two separate (2d) KDEs for the origin and destination cluster are fitted.

With this conditional handling the method can utilize good regional data availability but is still able to fallback into a mode with less freedom. Note that incomplete data which has been augmented with the procedure from Section S-I.1 is excluded from the fitting to keep the diffusion at a desired level.

S-I.3. Code for implementation of trip generation model

The model described in Section 3.2 uses several different techniques from ML and Statistics and needs to work with a large amount of data. Because of this requirement we built the framework with Python (version 3.7+), because many well suited and efficient libraries are available. To assist researchers who need to set up a similar framework, we included an overview of libraries used for this project:

- The relatively recent⁷ Python 3.8 is the programming language we choose for this project. To make calls to the dedicated OSRM-server as fast as possible, we are utilizing the python's asynchronous programming library `asyncio`, which was introduced in version 3.7.
- Numpy is the python library which offers fast numeric standard operations for arrays, e.g. matrix multiplication. Many other libraries listed here are build on numpy
- Scipy is build on Numpy and offers many function needed for scientific computing and basic machine learning. In this project we use the KDE implementation of `scipy`.
- SKLearn is a ML library containing fast implementation for many standard ML algorithms. From SKLearn we are using the GMM Implementation and the K-Means implementation (which is faster than the one of `scipy`).
- One major dependency of this project is the data science library Pandas, which is used to interact with data. Pandas implements many features known from the famous R-Language for statistics.

⁶The origin-destination dependence on cluster level is not affected.

⁷At the time of this writing Python 3.9 is the latest version of Python




Figure S3. The fitted GMMs (in red) for each week day (vertically) and each purpose (horizontally), the last column shows the joined GMMs for all purposes on a certain day. Note that every GMM is already joined for each origin-cluster and thus may contain more than three modes. In the most cases the matching is relatively accurate. The cases here showing bad results (e.g. work/education during the weekend) are in general unlikely to happen and therefore have only a minor overall impact.




Figure S4. Number of trips by mode in the dataset for the average weekday and weekend.

- Geopandas is an extension to pandas which comforts working with geo-spatial data. The library is primarily used for the conversion between the MiD-coordinate format and gps representation.
- Matplotlib (MPL) and the more recent seaborn is used for visualization purposes. Seaborn is build on top of MPL and pandas, offering improved graphics out of the box and in general better handling than MPL.
- For OSM-based routing, the standalone software OSRM is used. OSRM is able to handle thousands of routing request per second and thus well suited to estimate more realistic travel times and distances.
- The whole setup is wrapped with Docker/Docker-Compose. Docker wraps projects with containers, which can be seen as lightweight Virtual Machine (VM). Like this the project may set up on arbitrary systems without software-requirements except docker. Docker-compose is used to manage multiple containers simultaneously (The python project and OSRM is organized in different containers).
- Even though not used in this project, we want to pull attention to the recent library OSMNX, which combines pandas and the python network/graph library networkx for the use with OSM data. OSMNX might be used as a python alternative to OSRM here, to visualize road-networks, or for other graph theoretic applications.




Figure S5. Share of different trip-purposes. The purposes listed here differ between work, education, service and free time. In cases with sufficient data, a differentiation between the way towards and back from the destination is made.




Figure S6. Statistical dependence between different purposes and modality on the left and the strong dependence between distance and modality on the right.




Figure S7. Computation time for each day for different fleet sizes for the results in Table 2. The error bars indicate the minimum and maximum across the simulations, while the line tracks the median value. All simulations were run in Matlab on a Intel Core i9-9900K CPU @ 3.60GHz with 32 GB RAM.




Figure S8. Daily profits (excluding fixed costs) for different fleet sizes, prices, and costs. The three prices considered are: 0.5€/km (circle), 1€/km (x), 2€/km (triangle). The different line colors represent different cost of acquiring a vehicle (30,000€ to 100,000€) and variable cost per distance traveled (0.1€/km to 0.2€/km). All calculations assume a vehicle life of 5 years and exclude any other costs.