

Large-Scale Simulation of Shared Autonomous Vehicles: Integrating Supply and Demand Perspectives

Krishna Murthy Gurumurthy
Department of Civil, Architectural and Environmental Engineering
The University of Texas at Austin
gkmurthy10@utexas.edu

Felipe de Souza, Ph.D.*
Argonne National Laboratory, Energy Systems Division
9700 S. Cass Avenue
Argonne, IL 60439
fdesouza@anl.gov

Annesha Enam, Ph.D.
Argonne National Laboratory, Energy Systems Division
9700 S. Cass Avenue
Argonne, IL 60439
aenam@anl.gov

Joshua Auld, Ph.D.
Argonne National Laboratory, Energy Systems Division
9700 S. Cass Avenue
Argonne, IL 60439
jauld@anl.gov

January, 2020

Word Count: 7,097 words + 1 table x 250 words = 7,347 word equivalents (+ 6 figures)

Presented at the 99th Annual Meeting of the Transportation Research Board.

Under review for publication in the *Transportation Research Record*.

ABSTRACT

Transportation Network Companies (TNCs) have been steadily increasing the share of total trips in metropolitan areas across the world. Micro-modeling TNC operation is essential for large-scale transportation systems simulation. In this study, an agent-based approach for analyzing supply and demand aspects of ride-hailing operation is done using POLARIS, a high-performance simulation tool. On the demand side, a mode-choice model for the agent, and a vehicle-ownership model that informs this choice, is developed. On the supply side, TNC vehicle-assignment strategies, pickup and dropoff operations, and vehicle repositioning is modeled with congestion feedback, an outcome of the mesoscopic traffic simulation. Two case studies of Bloomington, Illinois and Chicago, Illinois are used to study the framework's computational speed for large-scale operations, and the impact of TNC fleets on a region's congestion patterns. Simulation results show that a zone-based vehicle-assignment strategy scales better than relying on matching closest vehicles to requests. For large regions like Chicago, large fleets are seen to be detrimental to congestion,

* Corresponding Author

especially in a future where more travelers are using TNCs. From an operational point of view, an efficient relocation strategy is critical for large regions with concentrated demand, but not regulating repositioning can worsen empty travel, and, consequently, congestion. The TNC simulation framework developed in this study is of special interest to cities and regions, since it can be used to model both demand and supply aspects for large regions at scale, and in reasonably-low computational time.

Keywords: Agent-based modeling, ride-hailing, shared autonomous vehicles, large-scale, integrated demand-and-supply modeling.

BACKGROUND

Transportation Network Companies (TNCs) are now an established, and convenient, transportation mode in metropolitan areas. Such services provide door-to-door transportation through traveler requests and traveler-driver interaction mediated by smartphone applications (1). The prevalence of these services will likely increase with the emergence of autonomous vehicles (2)(3)(4)(5), with TNCs serving riders with a fleet of shared autonomous vehicles (SAVs). Modeling and simulating such services is, therefore, essential for forecasting their impact on the transportation system in forthcoming years.

Recent research has explored the potential impacts of TNCs from an operations perspective, especially while considering an SAV fleet. Two studies, one with a hypothetical grid network (3) and another for Austin, Texas (6), found that one SAV can replace around 10 conventional vehicles with a reduction in total vehicle-miles traveled (VMT) with pooling, or dynamic ride-sharing (DRS). A Singapore study simulated taxi trips as SAVs and found that one-third of the existing fleet was sufficient to meet the demand (2). Rising VMT was found to be a concern, which is especially true with induced demand. A larger replacement rate was found in Stockholm, reaching 20 vehicles when DRS was employed (7). In Berlin, researchers simulated a fleet of 100,000 vehicles which was able to serve all trips previously performed by 1.1 million private cars (5). The same study also found an increase of 17% of the time spent in completing those trips. A Lisbon study simulated SAVs with different seating capacities to compare mode share changes (8). Using 4-, 8-, and 16-seaters, the authors observed the impact that solo travel and sharing rides had on the network when transit and walking modes were available. Large average vehicle occupancies (AVO) of 2.0 for 4-seaters, 4.2 for 8-seaters, and 11.4 for 16-seaters were accompanied with 20-40% VMT savings. The strategy to match vehicles to travelers was perceived as the key factor in reducing empty VMT (eVMT) and travelers' waiting time. Another study tested and compared several real-time dispatching strategies (9). The method was applied to the City of Chicago, USA, where taxi data is available. The optimization-based schemes fared better among all strategies, and the share of eVMT was about 20%. Similarly, the use of short-term demand forecasting with smart SAV relocation to meet anticipated demand has been studied (10). Their study helped show that even spatially aggregate forecasts of high quality help improve SAV performance. Two Austin-based studies pursued the impact of pricing SAVs and other modes with assumptions on rich information available from using AVs, concluding that for Austin, Texas, pricing with, and without, DRS helped curb the increase in VMT, but could not show considerable savings (11)(12). An interface has been created using the multi-agent simulation tool, or MATSim (15), to test SAV operations uniformly across case-studies (13)(14), with the only drawback of using a fixed demand and itinerary for every agent to begin with. Although the studies mentioned so far modeled SAV

operations in detail, the anticipated change in traveler behavior with SAVs was not incorporated in sync with SAV operations forecasting.

Travel behavior has typically been studied only with stated-preference surveys due to minimal data provided by TNCs, and the lack thereof for a non-existent future scenario. A survey with riders found important differences between taxi and ride-hailing travelers (16). Their results also showed that the increase in TNCs' mode share was occurring at the expense of other transportation modes. Key characteristics of SAV users were captured in another survey (17), finding that they were largely young, well-educated adults who had higher income, which was corroborated by (18) for TNC use. Higher ridership was observed in high-density areas as expected. Suburban population and older people were found to be less likely to adopt AVs (19), be it for ownership or for use in SAV fleets. Another survey studied the traveler behavior in terms of privacy and willingness to use DRS, and found that people are expected to be comfortable with DRS in the future (20). These aspects should be taken into account when forecasting future TNC ridership, in sync with SAV operational performance.

Many previous studies have highlighted the importance of integrated modeling in transportation (21)(22)(23)(24), but despite increasing literature in the area, very few studies have modeled TNC operation with explicit models for both demand and supply. The mobiTopp framework can be used to jointly model demand and supply aspects at the zone level (25). Their SAV simulation for Stuttgart, Germany, showed considerable VMT savings potential. Congestion effects on the links are not observed, and the SAV framework did not allow vehicle repositioning, which may reduce some of the savings but may be used by SAV fleet operators to maximize revenue. A study using MATSim incorporated an individual's willingness to use SAVs and user trust, and compared it with the in-house mode-choice model that uses a co-evolutionary algorithm and some randomness (26). This France-based study showed that incorporating the demand-side model considerably impacted the demographic for SAV use bringing more realism to the model. SimMobility is able to integrate supply and demand fully, and was used to compare current TNC operation with an SAV fleet for Singapore (27)(28), with results showcasing a benefit from centralized control. However, no study has been able to integrate supply and demand while being able to study a 24 hr simulation at scale (i.e., full population). The scale of analysis has especially been found to be of relevance for SAV operation when considering pooling (29).

This paper presents an agent-based approach for analyzing the operational aspects of TNCs at scale, while also using demand-side models to improve forecasting accuracy. The package is built on the top of the agent-based tool POLARIS (30) and integrates both demand and supply aspects. On the demand side, a vehicle ownership model is used at the household level to distribute vehicles based on household characteristics. A vehicle fleet distribution model takes care of the percentage of human-driven vehicles and AVs in the simulation. Cost of travel in a TNC vehicle and outcomes of the vehicle ownership and fleet distribution models inform the mode-choice model for a traveler, and this is repeated in a cycle for convergence. On the supply side, trip travel times and congestion effects are an outcome of the mesoscopic traffic assignment model, which also takes care of transit schedules and TNC vehicles for an integrated operations simulation. The implementation is generalized so that current TNC operation and future SAV systems can both be simulated.

The remaining sections of this paper are organized as follows: the agent-based model, POLARIS, and underlying principles are introduced next, followed by the methodology used to model TNCs

and SAVs, then concluding with the analysis conducted using two case studies, their results and discussion.

THE AGENT-BASED MODEL – POLARIS

The evolving complexity of travel behavior and transportation networks have led to the integrated modeling of travel demand and network supply as opposed to the traditional sequential process. In particular, agent-based models capture these interactions on the perspective of co-evolving agents that engage in activities (demand) based on time-dependent travel times using the different modes available (supply). The travel times, in turn, are an outcome of all the agents' decisions and the congestion effects in the network. Therefore, in these tools, demand and supply are mutually dependent as opposed to traditional four-step modeling. The sequential process practiced widely by many metropolitan planning organizations misses the interaction between demand generation and supply, and they will have to eventually shift to using tools like MATSim (15), SimMobility (31) or POLARIS (30). In this study, POLARIS was extended to micromodel TNC operations.

POLARIS is a high-performance, open-source, agent-based modeling framework that can simulate large-scale multi-modal transportation systems at a mesoscopic level that includes 100% of the population as agents in the model. It features integrated travel demand, network flow, and a traffic assignment model in which key aspects of travel decisions (activity planning, route choice, and tactical-level driving decisions) can be modeled simultaneously, and, in a continuous, fully integrated manner.

The activity models adapted from existing research controls each traveler's decision-making process for within-day, mid-term, and long-term timeframes taking into account activity types, and preferred modes and destinations (32)(33). The mid-term and within-day travel behavior decisions include the process of individual activity episode planning and engagement. These decisions are constrained by long-term choices regarding home and workplace choice, and household vehicle choices, and, in turn, these influence activity and travel planning, and realization.

The traffic flow model uses a mesoscopic representation of vehicles' movements based on Newell's kinematic wave model (34), with updates that represent interactions with traffic control infrastructure (35). The traffic flow simulation works seamlessly with a dynamic traffic assignment (DTA) algorithm (36) that assign routes to individual vehicles using a time-dependent A* shortest path router (37) based on the prevailing traffic condition, as well as updated skim travel times. Traveler's routing behavior from large delay is also captured by allowing re-routing. Effect of ITS infrastructure (30) and information mixing (36) on travel demand and behavior has also been incorporated, and the flow model takes into account connectivity and automation to observe likely behavior by personal AVs in the future (38). This functionality is inherited for the TNC model for futuristic SAV simulation. POLARIS also provides functionality for microscopic transit simulation. Transit lines and schedules in a region are simulated with capacity truncation in transit vehicles. Traveler and transit-vehicle interaction is captured for boarding and alighting, and trip delays take into account real-time transit movement based on congestion on the network through an iteratively-updated travel-time skim. The array of supply- and demand-side models in POLARIS, paired with low computational times, make it a suitable choice for large-scale TNC fleet simulation.

TNC FLEET SIMULATION

TNC fleets are modeled jointly from both the demand and supply perspectives in POLARIS. The demand side takes care of TNC request generation using a household vehicle ownership model and a mode-choice model that takes into account fares and travel times to identify the preferred mode. For the supply, TNC vehicle operation is modeled through two main components: the TNC fleet operator and the TNC vehicle. The demand and supply aspects of the TNC simulation used in this study are explained in detail next.

TNC Demand: Vehicle Ownership, Mode Choice and TNC Requests

Mode-Choice Model

POLARIS uses three separate mode-choice models depending on the activity purpose: home-based work/school, home-based other and non-home based, similar to traditional modeling. The nested-logit formulation used to model mode choice include nine modes: drive alone, TNC use, ride as passenger, walk, bike, bus with walk access, bus with drive access, rail with walk access, and rail with drive access. Among these, drive alone and TNC use are grouped under the auto nest, and the two rail modes are grouped under the rail nest, as shown in Figure 1. The model includes a variety of demographic variables, accessibility information, as well as level of service (LOS) variables. The demographic variables include individual demographics such as education, employment status, possession of driver’s license, and household demographics such as household income, household size, and vehicle and bike ownership among others. Road-network density and activity density of the destination zone are used to capture the characteristics of land use and the transportation network. The TNC-specific LOS variables used in the model include in-vehicle travel time and wait time (obtained from the simulation), and input fare. The TNC fare comprises of a fixed cost per trip, a distance-varying component, and a time-varying component. The model is developed and calibrated against the household travel survey data collected from the local region’s metropolitan planning organization.

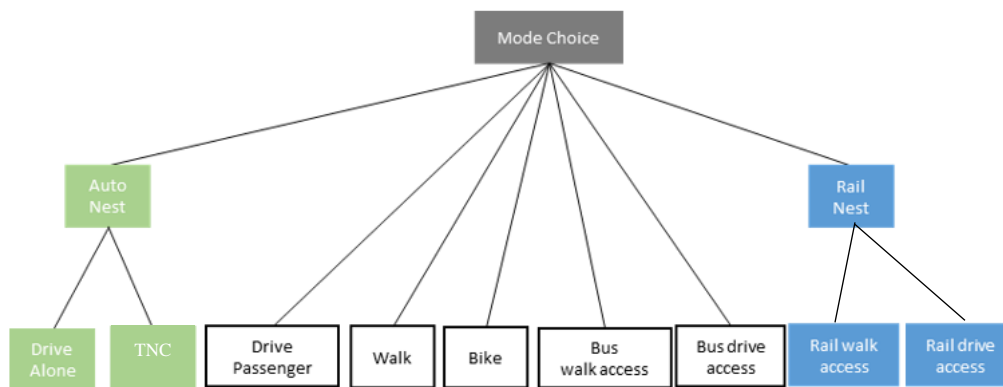


Figure 1 Nested-logit mode choice structure used in POLARIS

Vehicle-Ownership Model

Household vehicle ownership is expected to change in response to high penetration rate of TNCs. To account for the change in a household’s vehicle ownership, a household vehicle-disposal model

is implemented in POLARIS, in addition to the baseline vehicle ownership model. Two separate ordered-probit models are implemented to capture the vehicle-disposal propensity of one- and multi-vehicle households, respectively (39). The ordered-probit models quantify the household's likelihood of disposing vehicles from the existing household vehicle fleet. Consequently, the vehicle count is reduced for the households that express high likelihood of relinquishing their vehicle in response to increased penetration of TNCs. A random draw from a triangular distribution of unit base is used to determine the exact number of vehicles discarded by multi-vehicle households. The ordered-probit model uses a variety of demographic characteristics of the household (such as household size and income), as well as information of the primary driver in the household (such as gender, ethnicity, and education status). Travel characteristics such as the average travel duration on a regular day, and time spent in searching for parking are also included in the model. The model is sensitive to TNCs' service characteristics such as fare and fleet size. Accessibility of a household's location is also used to capture behavior leading to households located near urban core being more likely to give up vehicles compared to the households located in the suburbs, and is depicted in Figure 2 for Chicago, Illinois. The accessibility is calculated for each zone based on how attractive a destination zone is in terms of activity opportunities (e.g. government offices, service opportunities, manufacturing, retail and industrial opportunities) weighted by the travel time between the origin and destination zones. Activity opportunities at a zone farther away from the origin zone are weighted less favorably compared to the activity opportunities provided by a zone closer to the origin zone. Hence, in the simulation, a household having good (poor) accessibility score indicates that the household has good (poor) access to zones with favorable activity opportunities.

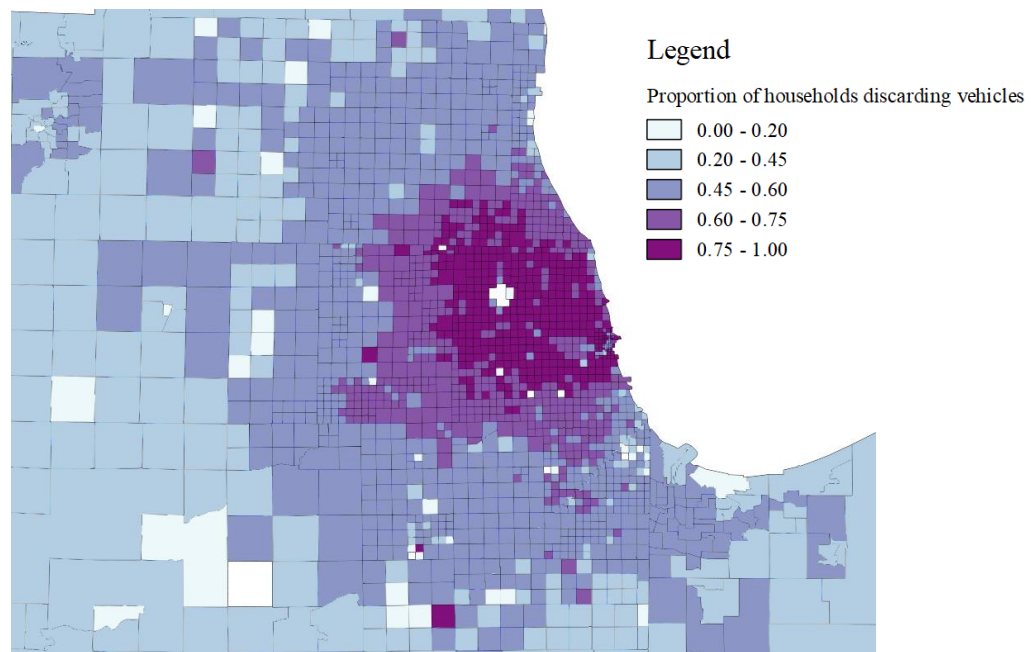


Figure 2 Distribution of households discarding vehicles in response to high TNC penetration in the Chicago region

TNC Supply

The TNC operator and TNC vehicle are modeled to mimic operations which are currently observed from a consumer's perspective. The operator assigns requests centrally to individual vehicles depending on the assignment strategy and monitors demand-to-supply ratio to determine if repositioning is required. The TNC vehicle executes the pickup, dropoff and repositioning tasks depending on the instruction received from the operator. The TNC vehicle is able to store requests that are being executed and those that need to be executed in the future.

Assignment Strategy

When a TNC trip request is made, the fleet operator attempts to assign it to the closest available vehicle in order to reduce total eVMT, as well as to minimize the traveler waiting time. Two assignment strategies are included in POLARIS: a coordinate-based search and a zone-based architecture. The coordinate-based search utilizes vehicle coordinates which are updated continuously throughout the simulation to create an R-tree (40). The R-tree improves spatial searches, and, therefore, the coordinates of an incoming request can be used to find the closest vehicle assessed using the Euclidian distance and the R-tree structure. Despite R-tree being an efficient spatial data structure, it may have some drawbacks. The Euclidian distance may not necessarily reflect the travel times between vehicles and requests since mean speeds vary along different routes in the network depending on congestion and network characteristics. Another aspect critical in large-scale simulations is the computational burden. R-tree's insertions, deletions, and queries take considerably more time compared to basic data structures (on the order of milliseconds compared to microseconds for the latter). However, this does not significantly grow with fleet size for individual operations. While this is not an issue to simulate present-day TNC systems that enjoy a relatively small share of trips, it may be significant for future scenarios in which millions of trips are potentially served by SAV fleets.

In order to account for this, an adapted version of the zone-based structure in (5) is implemented to be able to simulate large-scale scenarios with a reasonably-low computation time. The zone-based architecture is generated using the traffic analysis zones (TAZs) which is included in POLARIS. The TNC operator constructs an array of all neighboring zones, for each zone in the region, in ascending order of free-flow travel times defined with respect to a reference zone. This array is truncated using a predefined threshold for wait times so that a minimum level of service is maintained by the TNC service. By definition, the first zone in each array is the reference zone itself, which would mean minimal wait time. When a trip is requested, the operator checks for any available vehicle, starting from the origin zone and in the same order defined by the array, and assigns an unoccupied vehicle from the first neighboring zone that has a vehicle available. Since vehicles in these arrays are stored with longest-idling vehicle first, they need not be the closest available vehicle, but are expected to serve the request in reasonably less time.

Vehicle Operation

Once the request is assigned, the TNC vehicle handles the remainder of the request. Each vehicle stores a sorted list of tasks to be performed. A pickup and a dropoff operation is added to this list for every request assigned to the vehicle. Each task in the list involves the vehicle moving between its location to either the pickup point or the dropoff point. Depending on the task, the vehicle identifies the path from its current location to the next operation location. At the end of each trip, the total trip distance, travel time and empty travel, are computed and recorded. At the vehicle's destination after dropoff, the vehicle may receive a new set of tasks and repeat the same process again. If there is no task to be performed, the vehicle stays idle at the last task's destination

and awaits new trip assignment. Without DRS, only one task is executed at a time to speed up the assignment process, which only uses available vehicles (i.e., those that are idle).

Repositioning

Depending on the pattern of origins and destinations of incoming trip requests, there can be an uneven spatial distribution of the TNC fleet. The TAZs that have more trip beginnings tend to not have vehicles close to them to serve the incoming demand. Conversely, areas that are common destinations, but with a dearth in trip beginnings, have a concentration of vehicles. These vehicles will probably wait a long time to be assigned a new trip, and the consequent eVMT is likely to be long. Therefore, vehicles also need to be able to perform repositioning trips in addition to pickup and dropoff tasks.

The repositioning strategy in this study is undertaken in a decentralized manner, with vehicles performing repositioning trips based on their current location, and demand (number of requests), and supply (number of vehicles) aggregated at the TAZ level. The number of requests per TAZ is averaged over a pre-defined moving time-window, K . Based on the demand and the fleet size, S , a desirable number of vehicles per zone is computed as:

$$d_i = \frac{D_i + A_i}{\sum_j (D_j + A_j)} S,$$

where, d_i is the desired number of vehicles on zone i , D_i is the number of requests in the same zone, A_i is zone specific constant to a. Here, we set A_i to be proportional to the area of the zone, so that for low demands ($D_i \ll A_i$), the repositioning strategy attempts to distribute the vehicles evenly in the network. For high demands ($D_i \gg A_i$), the distribution will be proportional to the past demand. Similarly, the supply of a zone, s_i , is computed by counting the vehicles that is either idle in that zone i or serving requests in which the scheduled last dropoff is at zone i .

Based on the zone's desired number of vehicles, d_i , and supply, s_i , the repositioning decisions are undertaken. Vehicles in zones with ratio $r_i = s_i - 1/d_i > 1$, consider repositioning to zones in which the ratio r_i is lower than 1. To reduce computational time, a subset of potential destination zones are considered as opposed to verifying every potential destination. The specific zone that the vehicle repositions to is selected proportional to the following score:

$$u_i = \frac{1}{r_i t_i},$$

where, t_i is the predicted travel time from the vehicle current location to the destination. This probabilistic assumption is made to resemble the current TNC operation, and to serve as baseline for more sophisticated strategies in the future.

CASE STUDY AND RESULTS

The TNC model implemented in POLARIS comprises of several moving parts. There are two assignment strategies modeled, which can be compared for computational burden in large-scale simulations. Traveler interaction with TNC vehicles shows the fleet's level of service, and the TNC operation, of a TNC operator controlling a fleet of vehicles, can be studied to observe the impact on system VMT and induced congestion from eVMT. On the demand side, interest in TNC use based on vehicle-ownership and vehicle-disposal models and TNC use-costs helps understand the future of such a service.

The demand and supply perspectives implemented in this paper were tested using two case studies for Bloomington and Chicago in Illinois, United States. The two regions are a few orders of magnitude apart in terms of spatial extent, population, travel demand, and network complexity, making them good test beds. Bloomington houses about 120,000 people who make around 614,000 person-trips. On the other hand, the Chicago region is home to about 10 million people who make 27 million person-trips. Figure 3 shows the Bloomington network, which consists of about 4,000 links and 3,900 nodes, next to the Chicago network, which has about 31,000 links and 19,000 nodes.

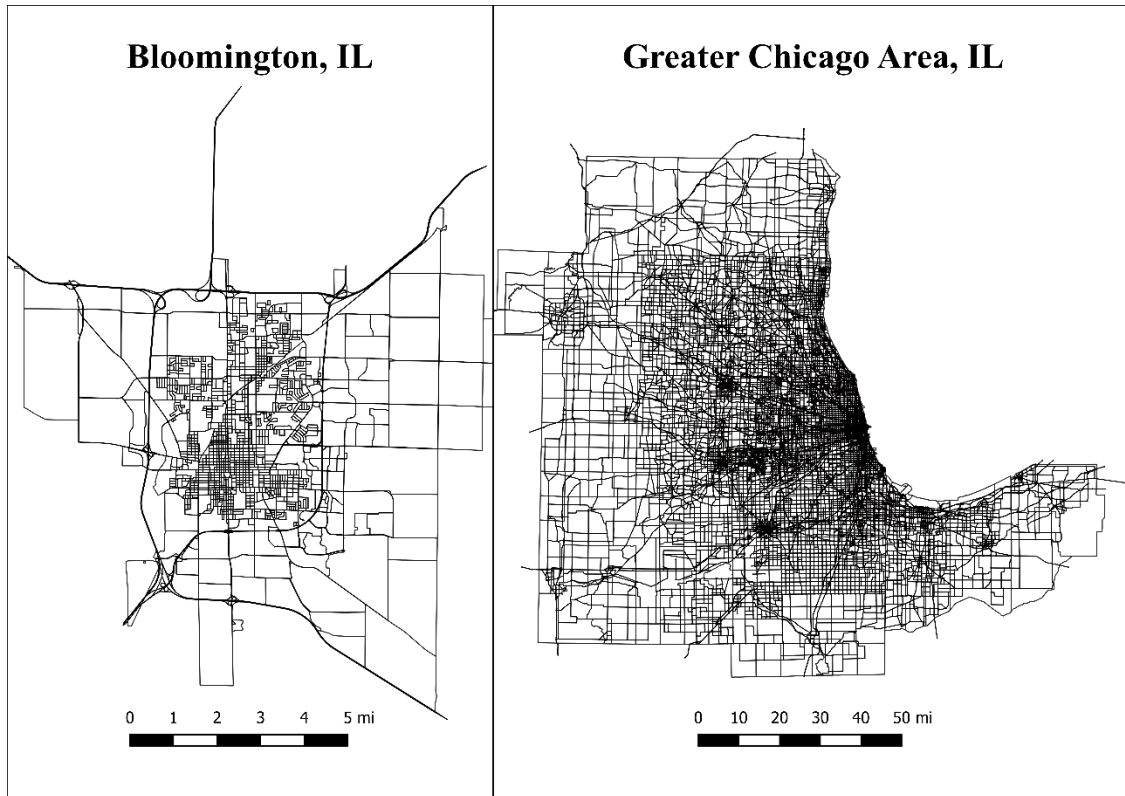


Figure 3 Bloomington and Chicago networks in Illinois, USA

Computational Burden and Effect on Wait Time

The smaller region of Bloomington was used in comparing the computational burden of TNC vehicle-to-request assignment for the two strategies that were modeled. The lower run time allowed simulation of multiple combinations of fleet sizes and maximum threshold wait times, which are important inputs to these strategies. In the present-day context, about 2-3% of all person-trips in Bloomington are made using TNCs and taxis. Fleet sizes varying as 1 TNC vehicle for every 75, 150, and 300 residents, were tested using the coordinate-based search and zone-based architecture, defined earlier. Since the zone-based architecture is created based on the threshold wait time allowed, wait times of 10, 15 and 20 min were simulated. The cumulative time taken in seconds for assignment is reported for each scenario, and the impact of these strategies on average wait time are discussed.

Figure 4 (a) illustrates the computational burden of each vehicle-assignment strategy implemented using simulations for Bloomington. R-tree has a significantly higher fixed computational cost,

about twice that of the zone-based architecture implemented. This only increases with fleet size although the average wait time is lowered by assigning closest vehicles to requests. The zone-based architecture performs equally well when taking into account the computational cost. Average wait times differ by less than a minute even when using this architecture as depicted in Figure 4 (b). The size of TAZs may play an important role in the effectiveness of the architecture, but was not analyzed since they were obtained from the local metropolitan planning organization's travel demand model.




Figure 4 Computation and traveler wait times for each vehicle-assignment strategy

Large-Scale Impacts from Repositioning

The larger region of Chicago was studied to validate the micro-modeling of TNC operations at scale. Fleet sizes of 1 TNC vehicle for every 25, 35, 50, and 100 residents, were tested for present-day conditions, and two levels of future reduction in vehicle ownership informed by the implemented model. The scenarios tested here are named A1-A4 for smaller change in vehicle ownership in the region, with 1-4 denoting increasing fleet size, or decreasing fleet availability. Similarly, B1-B4 denote the larger change in vehicle ownership that will considerably increase the demand for TNCs. The zone-based architecture was used for the assignment strategy with a maximum wait-time threshold of 30 min. Repositioning of TNC vehicles was allowed to occur at 10-min intervals based on the demand-supply balancing algorithm with the time-window of 4-hours for averaging trip requests. The average wait time, percent TNC trip requests served, percent eVMT and rVMT, and average trips made by each TNC vehicle is reported to analyze the impact of such fleets on the system.

Results from the large-scale simulation of the Chicago region were obtained from running POLARIS on a desktop with two Intel Xeon processors, each with 6 cores, running at 3.4GHz and having access to 128 GB of RAM. The convergent simulation of the Chicago travel demand model with TNC vehicles takes about 8 hours on the machine described with repositioning enabled. Table 1 summarizes the fleet characteristics for all the scenarios tested. In all the scenarios analyzed, there are trips unserved even when there is a higher vehicle availability. Most of the unserved requests occur during the morning peak as can be observed from Figure 5, which depicts the temporal distribution of the unserved requests across the different scenarios. In the morning peak, there is high demand of requests originating in the suburbs by travelers wanting to go to the central areas, causing an outflux of vehicles from that area. In addition to this, a reposition trip tends to be long as these areas are sparse. This explains the higher percent rVMT observed as fleet size increases across cases A and B.

The VMT in repositioning and pickup trips is a critical aspect of TNC/SAV operation as the vehicles travel empty (or without any customer), referred to as eVMT in this paper. Across scenarios A and B, the eVMT as a percentage of VMT stayed around 40%, with the exception of the base scenario. In the present-day context, TNC vehicles serving trips across the region were found to have about 41% eVMT by (41). As the repositioning strategy resembles the current TNC operation (with driver/vehicle driving to high demand areas without a central coordination), results in this study were close. In the base scenario, this share was higher as more vehicles were available to take long repositioning trips to suburban areas, which caused the base case to have the highest share of repositioning VMT (rVMT) across all scenarios. A more sophisticated repositioning and assignment strategy could potentially reduce both eVMT and rVMT. Additionally, there are maximum savings of about 2% in total vehicle hours traveled (VHT) in the A scenarios and a maximum of about 22% in the B scenarios. However, these savings may be higher if smaller fleets are more efficient. Large fleet sizes decrease saved travel times, especially when the mode share of single-occupant vehicles is high, like in scenarios A3 and A4.

Across both future scenarios of change in vehicle ownership, the TNC availability and percent rVMT was seen to be an influential factor in large percent eVMT values. Larger the TNC demand, the more imperative that regions regulate TNC fleet sizes. Small fleets are able to be more efficient with high average trips made per vehicle per day when there is higher demand, serving trips with similar average wait times. This becomes clear in Figure 6, which shows the share of vehicles performing different operations or is idling. Figure 6(a) shows the shares for cases A1 and A4, while Figure 6(b) shows the same for scenarios B1 and B4. For the smaller fleets (dashed lines, cases A1 and B1), most of the vehicles are either performing pickup or dropoff trips and the share of idle vehicles are small throughout the day. As the fleet increases, there are more idle vehicles and some of these vehicles are able to perform repositioning trips to other areas of the network.

Table 1 TNC Fleet Impacts by Size and Demand

Sc.	TNC Availability (1 for x residents)	#TNC Trips per 100 residents	% eVMT (% of SAV VMT)	% rVMT (% of SAV VMT)	Avg. Trips per TNC vehicle per day	Avg. Wait Time (in min)	% Served Trips
Base	35	10 trips	49.0	36.0	6.7	7.9	97.1
A1	100	38 trips	43.0	10.0	26.3	13.4	70.5
A2	50		41.4	19.3	17.1	9.8	89.6
A3	35		42.3	24.6	12.7	8.5	93.7
A4	25		44.6	29.1	10.3	8.1	95.6
B1	100	65 trips	44.2	8.1	32.1	13.2	50.1
B2	50		40.6	11.0	25.6	11.9	79.4
B3	35		39.8	17.2	19.7	9.3	90.0
B4	25		40.9	21.7	15.5	8.4	93.1

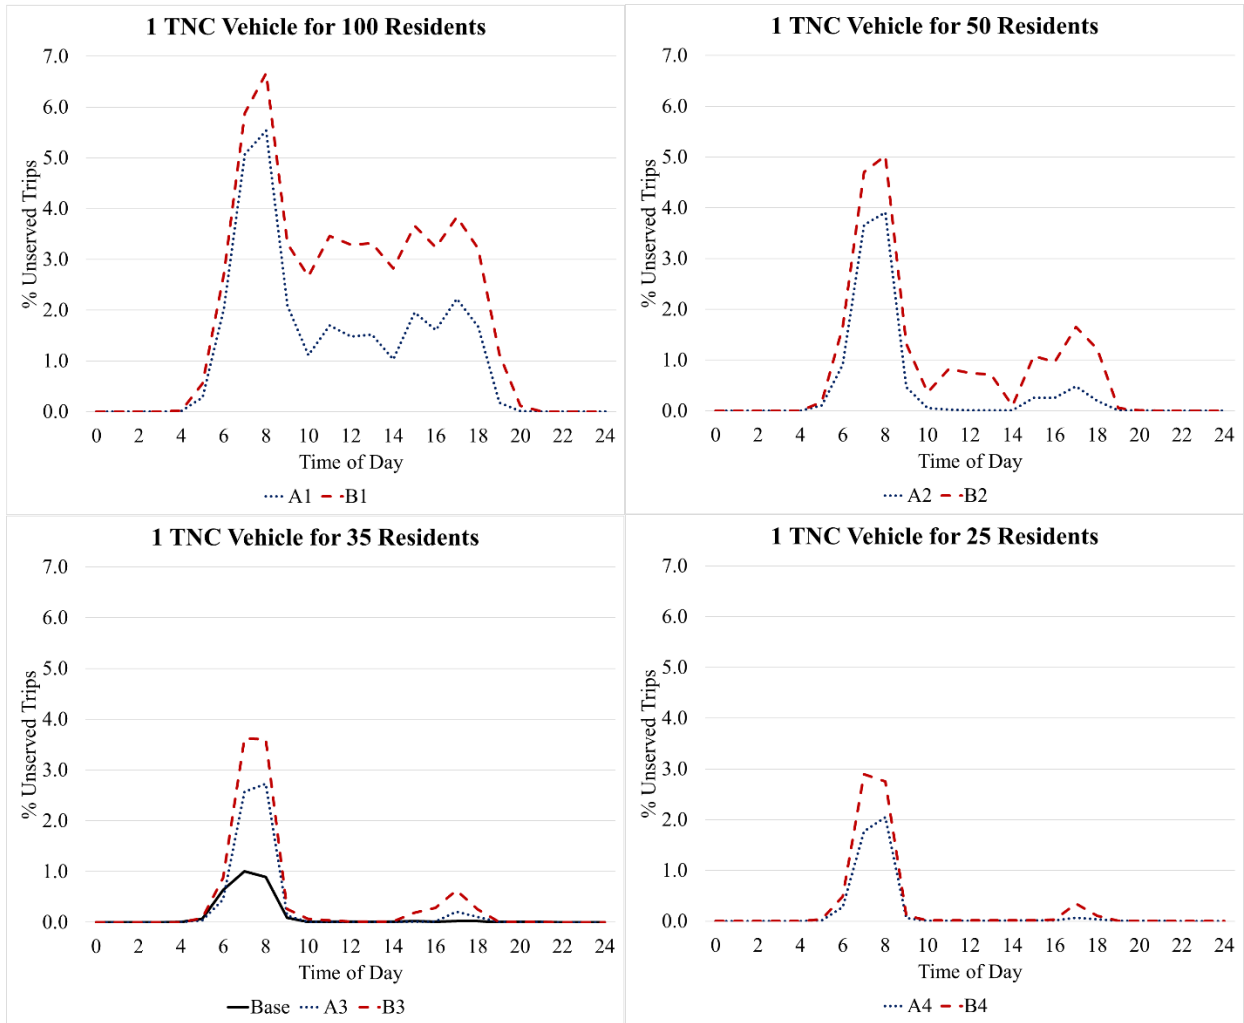


Figure 5 Percent requests unserved by time-of-day across scenarios

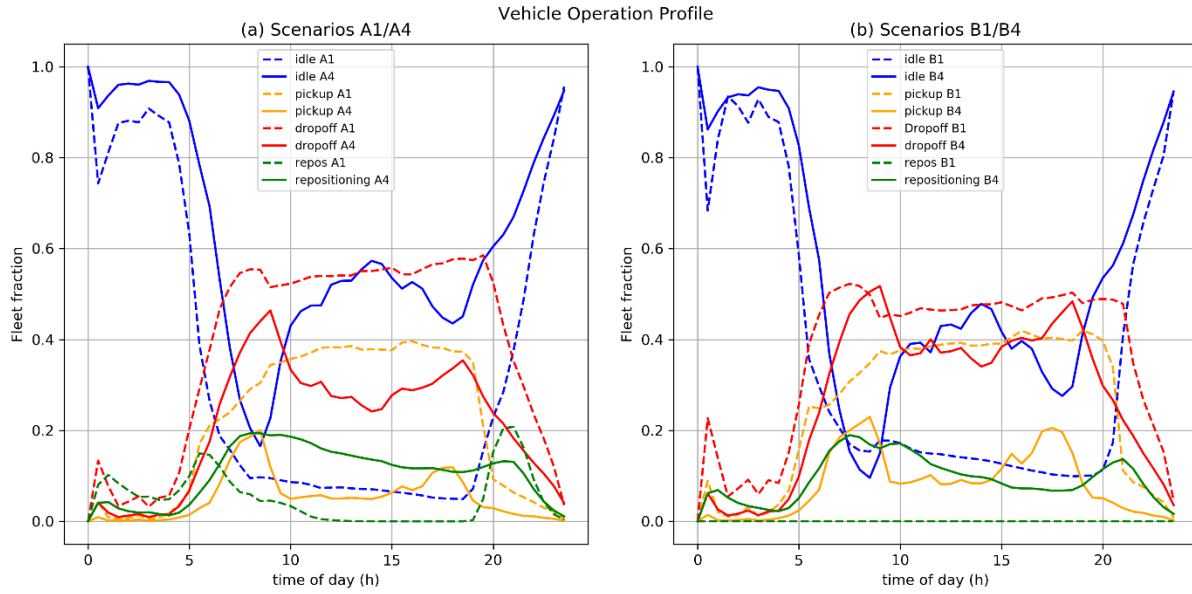


Figure 6 Share of vehicles performing pickup, dropoff, repositioning trips or idle by time of day for scenarios (a) A1 and A4, (b) B1 and B4.

CONCLUSIONS

Ride-hailing use has seen exceptional growth, and this is only expected to increase exponentially with AVs in the mix. TNC operation is implemented in POLARIS in this study, to simulate the interaction between TNC vehicles, travelers, and the network, accurately, and to allow for different operational strategies. This supply model is integrated with the demand side by modeling TNC requests at the household level depending on vehicle ownership and a mode-choice model. The framework supports simulation of large regions at scale in reasonably low computational time.

A case study for Bloomington, Illinois is used to show that a heuristic vehicle-to-request assignment strategy compares well with the closest vehicle search using a specialized data structure. The zone-based architecture is faster, and insensitive to the fleet size at the expense of a marginal increase in pickup distance and pickup travel time. Overall, both strategies lead to similar wait times, depending only on higher TNC vehicle availability to lower traveler wait times. Despite being effective, the closest vehicle search would significantly increase the computational time with millions of trips being assigned in quick succession. The zone-based implementation is then validated for large-scale use by simulating travel in Chicago, Illinois. Results from this analysis reveal that repositioning may be the reason behind lower eVMT reported in the literature. Although the algorithm to balance demand and supply implemented here is a heuristic, fleet sizes need to be regulated for anticipated changes in TNC demand so as to keep added VMT in check. This anticipated change in demand, replicated here using a vehicle disposal model informs regions regarding the concentration of new demand in urban cores, where households are likely to own fewer vehicles. Percent eVMT was around 40% across all scenarios, and the percent rVMT increases as the fleet size increases, irrespective of demand. A small, but non-zero, percentage of TNC demand was not served even when employing larger fleet sizes. The unserved trips are concentrated in the morning peaks, and in suburban areas.

The methodology here is a step forward in integrating demand and supply for large regions. There were some key drawbacks that were observed during the analysis. On the demand side, TNCs operate based on drivers for-hire, serving as contractors, who drop in and dropout of working for the TNC during different times of the day. A constant fleet assumption modeled here can be improved using time-dependent fleet sizing based on dynamic pricing. Similarly, a thorough TNC demand model needs to include human tendency for cruising similar to (28). Once pooling is modeled, a choice dimension for choosing to pool can also improve realism and is ongoing work. On the supply side, the zone-based architecture implemented relies on the underlying framework of TAZs that is typically used by local metropolitan organizations. They are not uniformly distributed across a region, and this can play an important role in the effectiveness of this assignment strategy. The repositioning of trips using a demand-and-supply balancing algorithm is a heuristic that does not take into account a maximum distance or travel time that a repositioning trip should last. When this is combined with an ineffective TAZ size, it can lead to larger than usual levels of repositioning. DRS options that are offered by TNCs, nowadays, can save travelers some money and help lower congestion, but was not explored in this study. Similar evaluations, as done in this study, needs to be performed along with modeling DRS. Results originating from using heuristics that are used here needs to be compared with optimization-based strategies for vehicle-assignment and repositioning for a TNC/SAV fleet operation as it can significantly impact eVMT and VHT. These are out of the scope of this study, but are in the works.

ACKNOWLEDGEMENTS

This report and the work described were sponsored by the U.S. Department of Energy Vehicle Technologies Office under the Systems and Modeling for Accelerated Research in Transportation Mobility Laboratory Consortium, an initiative of the Energy Efficient Mobility Systems Program. David Anderson, a Department of Energy Office of Energy Efficiency and Renewable Energy manager, played an important role in establishing the project concept, advancing implementation, and providing ongoing guidance.

AUTHOR CONTRIBUTION STATEMENT

The authors confirm contribution to the paper as follows: study conception and design: Gurumurthy, K.M., de Souza, F., and Enam, A.; analysis and interpretation of results: Gurumurthy, K.M., and de Souza, F.; draft manuscript preparation: Gurumurthy, K.M., de Souza, F., Enam, A., and Auld, J. All authors reviewed the results and approved the final version of the manuscript.

REFERENCES

1. Shaheen, S. A., A. Coehn, I. Zohdy, and B. Kock. *Smartphone Applications to Influence Travel Choices: Practices and Policies*. Publication FHWA-HOP-16-023. FHWA, 2016.
2. Spieser, K., K. Treleaven, R. Zhang, E. Frazzoli, D. Morton, and M. Pavone. Toward a Systematic Approach to the Design and Evaluation of Automated Mobility-on-Demand Systems: A Case Study in Singapore. In *Road Vehicle Automation* (G. Meyer and S. Beiker, eds.), Springer International Publishing, Cham, pp. 229–245.
3. Fagnant, D. J., and K. M. Kockelman. The Travel and Environmental Implications of Shared Autonomous Vehicles, Using Agent-Based Model Scenarios. *Transportation Research Part C: Emerging Technologies*, Vol. 40, 2014, pp. 1–13. <https://doi.org/10.1016/j.trc.2013.12.001>.

4. Fagnant, D. J., and K. Kockelman. Preparing a Nation for Autonomous Vehicles: Opportunities, Barriers and Policy Recommendations. *Transportation Research Part A: Policy and Practice*, Vol. 77, 2015, pp. 167–181. <https://doi.org/10.1016/j.tra.2015.04.003>.
5. Bischoff, J., and M. Maciejewski. Simulation of City-Wide Replacement of Private Cars with Autonomous Taxis in Berlin. *Procedia Computer Science*, Vol. 83, 2016, pp. 237–244. <https://doi.org/10.1016/j.procs.2016.04.121>.
6. Fagnant, D. J., and K. M. Kockelman. Dynamic Ride-Sharing and Fleet Sizing for a System of Shared Autonomous Vehicles in Austin, Texas. *Transportation*, Vol. 45, No. 1, 2018, pp. 143–158. <https://doi.org/10.1007/s11116-016-9729-z>.
7. Burghout, W., P.-J. Rigole, and I. Andreasson. Impacts of Shared Autonomous Taxis in a Metropolitan Area. Presented at the Transportation Research Board 94th Annual Meeting Transportation Research Board, 2015.
8. Martinez, L. M., and J. M. Viegas. Assessing the Impacts of Deploying a Shared Self-Driving Urban Mobility System: An Agent-Based Model Applied to the City of Lisbon, Portugal. *International Journal of Transportation Science and Technology*, Vol. 6, No. 1, 2017, pp. 13–27. <https://doi.org/10.1016/j.ijtst.2017.05.005>.
9. Hyland, M., and H. S. Mahmassani. Dynamic Autonomous Vehicle Fleet Operations: Optimization-Based Strategies to Assign AVs to Immediate Traveler Demand Requests. *Transportation Research Part C: Emerging Technologies*, Vol. 92, 2018, pp. 278–297. <https://doi.org/10.1016/j.trc.2018.05.003>.
10. Dandl, F., M. Hyland, K. Bogenberger, and H. S. Mahmassani. Evaluating the Impact of Spatio-Temporal Demand Forecast Aggregation on the Operational Performance of Shared Autonomous Mobility Fleets. *Transportation*, 2019. <https://doi.org/10.1007/s11116-019-10007-9>.
11. Simoni, M. D., K. M. Kockelman, K. M. Gurumurthy, and J. Bischoff. Congestion Pricing in a World of Self-Driving Vehicles: An Analysis of Different Strategies in Alternative Future Scenarios. *Transportation Research Part C: Emerging Technologies*, Vol. 98, 2019, pp. 167–185. <https://doi.org/10.1016/j.trc.2018.11.002>.
12. Gurumurthy, K. M., K. M. Kockelman, and M. D. Simoni. Benefits and Costs of Ride-Sharing in Shared Automated Vehicles across Austin, Texas: Opportunities for Congestion Pricing. *Transportation Research Record*, Vol. 2673, No. 6, 2019, pp. 548–556. <https://doi.org/10.1177/0361198119850785>.
13. Hörl, S., C. Ruch, F. Becker, E. Frazzoli, and K. W. Axhausen. Fleet Control Algorithms for Automated Mobility: A Simulation Assessment for Zurich. Presented at the 97th Annual Meeting of the Transportation Research Board, 2018.
14. Ruch, C., S. Hörl, and E. Frazzoli. AMoDeus, a Simulation-Based Testbed for Autonomous Mobility-on-Demand Systems. Presented at the 2018 21st International Conference on Intelligent Transportation Systems (ITSC), 2018.
15. Horni, A., K. Nagel, and K. W. Axhausen, Eds. *The Multi-Agent Transport Simulation MATSim*. Ubiquity Press, 2016.
16. Rayle, L., D. Dai, N. Chan, R. Cervero, and S. Shaheen. Just a Better Taxi? A Survey-Based Comparison of Taxis, Transit, and Ridesourcing Services in San Francisco. *Transport Policy*, Vol. 45, 2016, pp. 168–178. <https://doi.org/10.1016/j.tranpol.2015.10.004>.
17. Krueger, R., T. H. Rashidi, and J. M. Rose. Preferences for Shared Autonomous Vehicles. *Transportation Research Part C: Emerging Technologies*, Vol. 69, 2016, pp. 343–355. <https://doi.org/10.1016/j.trc.2016.06.015>.

18. Dias, F. F., P. S. Lavieri, V. M. Garikapati, S. Astroza, R. M. Pendyala, and C. R. Bhat. A Behavioral Choice Model of the Use of Car-Sharing and Ride-Sourcing Services. *Transportation*, Vol. 44, No. 6, 2017, pp. 1307–1323. <https://doi.org/10.1007/s11116-017-9797-8>.
19. Lavieri, P. S., V. M. Garikapati, C. R. Bhat, R. M. Pendyala, S. Astroza, and F. F. Dias. Modeling Individual Preferences for Ownership and Sharing of Autonomous Vehicle Technologies. *Transportation Research Record*, Vol. 2665, No. 1, 2017, pp. 1–10. <https://doi.org/10.3141/2665-01>.
20. Gurumurthy, K. M., and K. M. Kockelman. Modeling Americans' Autonomous Vehicle Preferences: A Focus on Dynamic Ride-Sharing, Privacy & Long-Distance Mode Choices. *Technological Forecasting and Social Change*, Vol. 150, 2020, p. 119792. <https://doi.org/10.1016/j.techfore.2019.119792>.
21. Hao, J. Y., M. Hatzopoulou, and E. J. Miller. Integrating an Activity-Based Travel Demand Model with Dynamic Traffic Assignment and Emission Models: Implementation in the Greater Toronto, Canada, Area. *Transportation Research Record*, 2010. <https://doi.org/10.3141/2176-01>.
22. Pendyala, R. M., K. C. Konduri, Y.-C. Chiu, M. Hickman, H. Noh, P. Waddell, L. Wang, D. You, and B. Gardner. Integrated Land Use–Transport Model System with Dynamic Time-Dependent Activity–Travel Microsimulation: *Transportation Research Record*, 2012. <https://doi.org/10.3141/2303-03>.
23. Rasouli, S., and H. Timmermans. Activity-Based Models of Travel Demand: Promises, Progress and Prospects. *International Journal of Urban Sciences*, Vol. 18, No. 1, 2014, pp. 31–60. <https://doi.org/10.1080/12265934.2013.835118>.
24. Lu, Y., K. Basak, C. Carrion, H. Loganathan, M. Adnan, F. C. Pereira, V. H. Saber, and M. Ben-Akiva. SimMobility Mid-Term Simulator: A State of the Art Integrated Agent Based Demand and Supply Model. Presented at the Transportation Research Board 94th Annual Meeting Transportation Research Board, 2015.
25. Heilig, M., T. Hilgert, N. Mallig, M. Kagerbauer, and P. Vortisch. Potentials of Autonomous Vehicles in a Changing Private Transportation System – a Case Study in the Stuttgart Region. *Transportation Research Procedia*, Vol. 26, 2017, pp. 13–21. <https://doi.org/10.1016/j.trpro.2017.07.004>.
26. Vosooghi, R., J. Kamel, J. Puchinger, V. Leblond, and M. Jankovic. Robo-Taxi Service Fleet Sizing: Assessing the Impact of User Trust and Willingness-to-Use. *Transportation*, 2019. <https://doi.org/10.1007/s11116-019-10013-x>.
27. Basu, R., A. Araldo, A. P. Akkinapally, B. H. Nahmias Biran, K. Basak, R. Seshadri, N. Deshmukh, N. Kumar, C. L. Azevedo, and M. Ben-Akiva. Automated Mobility-on-Demand vs. Mass Transit: A Multi-Modal Activity-Driven Agent-Based Simulation Approach. *Transportation Research Record*, Vol. 2672, No. 8, 2018, pp. 608–618. <https://doi.org/10.1177/0361198118758630>.
28. Nahmias-Biran, B., J. B. Oke, N. Kumar, K. Basak, A. Araldo, R. Seshadri, A. Akkinapally, C. L. Azevedo, and M. Ben-Akiva. From Traditional to Automated Mobility on Demand: A Comprehensive Framework for Modeling On-Demand Services in SimMobility: *Transportation Research Record*, 2019. <https://doi.org/10.1177/0361198119853553>.
29. Yan, H., K. M. Kockelman, and K. M. Gurumurthy. Agent-Based Simulation for Shared Autonomous Vehicle Use Across the Minneapolis-Saint Paul Region. , 2019.

30. Auld, J., M. Hope, H. Ley, V. Sokolov, B. Xu, and K. Zhang. POLARIS: Agent-Based Modeling Framework Development and Implementation for Integrated Travel Demand and Network and Operations Simulations. *Transportation Research Part C: Emerging Technologies*, Vol. 64, 2016, pp. 101–116. <https://doi.org/10.1016/j.trc.2015.07.017>.
31. Adnan, M., F. C. Pereira, C. L. Azevedo, K. Basak, M. Lovric, S. Raveau, Y. Zhu, J. Ferreira, C. Zegras, and M. Ben-Akiva. SimMobility: A Multi-Scale Integrated Agent-Based Simulation Platform. Presented at the Transportation Research Board 95th Annual Meeting Transportation Research Board, 2016.
32. Auld, J., and A. Mohammadian. Framework for the Development of the Agent-Based Dynamic Activity Planning and Travel Scheduling (ADAPTS) Model. *Transportation Letters*, Vol. 1, No. 3, 2009, pp. 245–255. <https://doi.org/10.3328/TL.2009.01.03.245-255>.
33. Auld, J., and A. Mohammadian. Activity Planning Processes in the Agent-Based Dynamic Activity Planning and Travel Scheduling (ADAPTS) Model. *Transportation Research Part A: Policy and Practice*, Vol. 46, No. 8, 2012, pp. 1386–1403. <https://doi.org/10.1016/j.tra.2012.05.017>.
34. Newell, G. F. A Simplified Theory of Kinematic Waves in Highway Traffic, Part I: General Theory - ScienceDirect. Vol. 27, No. 4, 1993, pp. 281–287.
35. Souza, F. de, O. Verbas, and J. Auld. Mesoscopic Traffic Flow Model for Agent-Based Simulation. *Procedia Computer Science*, Vol. 151, 2019, pp. 858–863. <https://doi.org/10.1016/j.procs.2019.04.118>.
36. Auld, J., O. Verbas, and M. Stinson. Agent-Based Dynamic Traffic Assignment with Information Mixing. *Procedia Computer Science*, Vol. 151, 2019, pp. 864–869. <https://doi.org/10.1016/j.procs.2019.04.119>.
37. Verbas, Ö., J. Auld, H. Ley, R. Weimer, and S. Driscoll. Time-Dependent Intermodal A* Algorithm: Methodology and Implementation on a Large-Scale Network. *Transportation Research Record*, Vol. 2672, No. 47, 2018, pp. 219–230. <https://doi.org/10.1177/0361198118796402>.
38. Auld, J., O. Verbas, M. Javanmardi, and A. Rousseau. Impact of Privately-Owned Level 4 CAV Technologies on Travel Demand and Energy. *Procedia Computer Science*, Vol. 130, 2018, pp. 914–919. <https://doi.org/10.1016/j.procs.2018.04.089>.
39. Menon, N., N. Barbour, Y. Zhang, A. R. Pinjari, and F. Mannering. Shared Autonomous Vehicles and Their Potential Impacts on Household Vehicle Ownership: An Exploratory Empirical Assessment. *International Journal of Sustainable Transportation*, Vol. 13, No. 2, 2019, pp. 111–122. <https://doi.org/10.1080/15568318.2018.1443178>.
40. Guttman, A. R-Trees. 1984.
41. Henao, A., and W. E. Marshall. The Impact of Ride Hailing on Parking (and Vice Versa). *Journal of Transport and Land Use*, Vol. 12, No. 1, 2019. <https://doi.org/10.5198/jtlu.2019.1392>.