CURB ALLOCATION AND PICK-UP DROP-OFF AGGREGATION FOR A SHARED AUTONOMOUS VEHICLE NETWORK

Christian B. Hunter
Graduate Research Assistant
Department of Civil, Architectural and Environmental Engineering
The University of Texas at Austin
christianhunter@utexas.edu

Kara M. Kockelman
Dewitt Greer Centennial Professor in Transportation Engineering
Department of Civil, Architectural and Environmental Engineering
The University of Texas at Austin – 6.9 E. Cockrell Jr. Hall
Austin, TX, 78712-1076
kkockelm@mail.utexas.edu

Shadi Djavadian
New Mobility Modeling Engineer
Robotics & Mobility Research Department
Ford Motor Company
Ford Greenfield Labs, Palo Alto, CA, 94304
sdjavadi@ford.com

Accepted for presentation at the 101st Annual Meeting of the Transportation Research Board

ABSTRACT

Advances in communication and information technologies and automation in vehicles have led to the birth of new transportation services such as shared autonomous vehicles (SAVs). SAVs are on-demand services with flexible routes and schedules, which can replace personal vehicles for many trip types in the near future, thereby reducing the number of personal vehicles on the road at rush hour. Pick-up and drop-off (PUDO) locations and densities are pressing questions for planning fleet operations of SAVs since they impact SAV demand, operation, and traffic congestion. Unlike traditional human-driven taxis and ride-hailing vehicles, SAVs cannot nimbly maneuver to access the curb or driveways, nor can they engage in quasi-legal procedures like double parking or fire hydrant pick-ups. As such, in order to safely and efficiently operate the vehicles they must have clear curb space to make their pick-ups and drop-offs of passengers. This study models the impact of different PUDO locations and densities on the fleet operation of SAV, potential SAV demand, and traffic congestion. The City of Austin, Texas, is used as a case study modeled in POLARIS, an agent-based simulation platform developed by Argonne National Laboratory. The results show that, SAV ridership increases with fleet size, especially during peak hours, but suffers from diminishing returns. For a network with 4000 SAVs that charges $1/mile, placing PUDOs every three blocks instead of every block increases the number of spaces needed at each PUDO site, reduces the total number of spaces required by 40%, and decreases SAV ridership by 6%.

BACKGROUND

Cities are facing increasing demands for curb access from diverse users such as on-demand micromobility, transportation network companies (TNCs), and urban freight delivery such that dedicating curb space for parking in dense urban centers is becoming less tenable (1). Agencies in charge of parking management have responded to new users by increasing staffing for greater enforcement and piloting new mobility zones, particularly for TNCs in nightlife and entertainment districts (2). A review of pick-up and drop-off zones (PUDOs) for TNCs estimates that they reduce operational failures arising from the demand for curb space exceeding capacity, resolving issues in pedestrian safety, traffic congestion, conflicts with bike lane users, and impairment of emergency vehicles (3) . The transition to a sharing economy with the expected arrival of shared autonomous vehicles (SAVs), which do not require a driver and carry dozens of different users per day, warrants further exploratory analysis on the operational performance of PUDOs.

The demand for shared mobility is not fixed and can vary from one day to another depending on several factors. One of the crucial factors affecting the choice of travelers to use shared mobility is PUDO location which can affect their wait time along with their access/egress time. In addition, PUDO locations can impact level of service of SAV (e.g. response time) and fleet size, which in turn can also impact the choice of travelers. As such it is important to capture the interaction between demand and supply and impact of PUDO location on mode choice (demand for SAV) of travelers and SAV fleet operations. The aim of this study is to simulate a network of SAVs in Austin, Texas in POLARIS and evaluate the effect of different PUDO location and density configurations on various aspects of performance such as wait time, walk time, average vehicle occupancy, and vehicle miles traveled in an integrated supply and demand context. . By varying the presence and spacing of PUDOs (taking into account curb access limitations and network link and zone exclusions) and SAV fleet size, an evaluation could be made of the impact of these variables and provide a valuable resource for municipalities considering how to accommodate an increase in curb demand due to coming SAVs. It can also provide potential SAV operators a window into the performance of an example network and allow them to show local governments their impact on street and/or curb congestion or the lack thereof, and lobby for dedicated PUDOs if necessary.

The research questions this paper attempts to address include:

- How does PUDO density impact the number of trips served?
- How do different PUDO characteristics impact SAV fleet performance (e.g. trips served per vehicle, deadhead miles, etc.)?
- What demands do SAVs place on curb space, and what is the appropriate number of spots at each PUDO location?

The reminder of this paper is organized as follows. In the literature review section, an overview of existing literature on SAVs modeling and evaluation is presented. The methodology section presents PUDO locations evaluation experiment design and simulation modeling. Then the case study of Austin, Texas is presented followed by results, analysis, summary, and future work.

LITERATURE REVIEW

SAVs differ from personal autonomous vehicles in that they are not owned by an individual but rather by a fleet operator. Instead of sitting in a parking lot once a trip is completed, they may drive themselves to begin another person's trip. In addition, SAVs can carry unrelated riders who share similar destinations. They also differ from current TNCs because they do not require a driver, which removes some of the cost of operation.

Vehicle dispatching has all been separately studied across many contexts, with conclusions that are relevant to SAV research. A thorough examination of event-based logic for dispatching shared autonomous vehicles in an agent-based simulation with congestion feedback is found in Levin et al. (4)**.** Bösch et al. (5) simulates low automated vehicles (AV) penetration in Zurich to estimate fleet sizes required to meet different demand levels, though all demand levels are less than ten percent of the total travel demand in the city. Fagnant et al. (6) described algorithms for distributed ride sharing including heuristics for reallocating idle SAVs that are frequently cited. Both are accomplished with some abstraction of the network as well as trip sampling. Hörl (7) simulated an autonomous taxi service with dynamic demand response, demonstrated on a toy model loosely based on Sioux Falls. Bischoff and Maciejewski (8) simulated AVs in Berlin using MATSim, but only 10% of trips are simulated because of computational limitations. For assignment, the authors use the common heuristic shortcut of assigning the closest vehicle to a request, and because they are not simulating shared vehicles, they avoid the complexity of the multiple vehicle pickup and delivery problem (MVPDP).

Loeb and Kockelman (9) used a logit model for wait time trip rejection when modeling shared, autonomous, electric vehicles (SAEVs) in Austin, Texas and examined how the fleet performs under several parameters including charge time and fleet size. Focus was given to the costs associated with SAEV fleets and infrastructure. However, the model did not include mode choice beyond rejecting trips, opting instead to fix SAEV demand at various levels and examine the costs of operating and charging the fleet. This paper follows Kockelman et al. (10) in placing a high volume of charging stations and then trimming based on utilization.

Curb Allocation

Predictions of the effects AVs on parking spaces and curb usage depend on scenario formulations of future vehicle ownership (privately-owned or fleet operator-owner), AV saturation rate, and the share of people willing to share rides with strangers (11). In a world of well-received SAVs, high utilization rates and vehicle occupancies in a near continuous operation could eliminate the need for off-street parking. In the short term, lessons taken from managing curb space arising from TNCs and the emergent SAVs will inform policies to manage traffic at inter-/multi-modal transportation facilities and major destinations (e.g., office parks, universities, concert venues, nightlife districts). Cities should anticipate a large demand for passenger loading/ unloading at destinations of TNC and future SAV trips, and manage on-street parking to increase turnover or designated PUDO locations to meet demand (1).

PUDO Modeling Techniques

In order to accommodate passengers beginning and ending their trips at different locations without adding to the trip length of all users of the shared vehicle, riders can be picked up and dropped off at designated locations. Riders walk to or from these sites instead of forcing the vehicle to make the trip all the way to each trip end (12). When combined with dynamic ridesharing (DRS), which enables SAVs that are already carrying passengers to pick up additional riders who are not far off the SAV’s current path, greater vehicle occupancy and or lower fleet sizes could theoretically be achieved. An example of DRS and PUDOs leading to reduced SAV trip length is shown below in Figure 1.

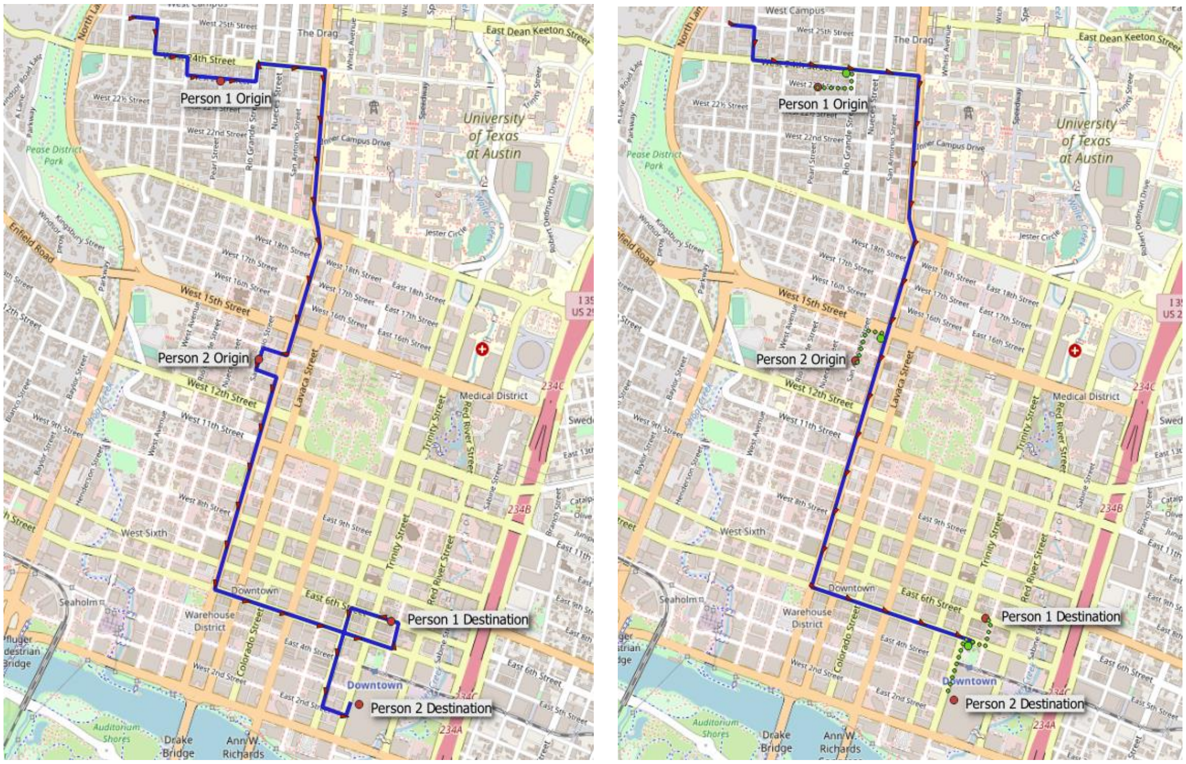

(a) (b)

FIGURE 1 DRS without (a) and with (b) PUDOs

In Figure 1(a), where there are no PUDOs, an SAV fulfilling two trips near the Austin, Texas CDB must leave the arterial it is traveling on several times to pick up a passenger at their exact origin point (see blue line), as well as take a convoluted path along several one-way streets at the end of the journey. With PUDOs, the two riders walk to meet the SAV (path shown in green dots) without it having to make a detour onto side streets, then both are dropped off at the same PUDO and walk a few blocks to their final destination. This yields a reduced trip distance and reduced wait time for all involved.

Modeling PUDO stop aggregation for SAVs is more complex than matching a single user to an SAV vehicle which then picks up the passenger. Rather, Fielbaum et al. (13) suggest one must consider vehicles passing near the origin and choose not only which vehicle to assign but also to what point to send the user on foot to meet the SAV.

They expand on the framework created by Alonso-Mora et al. (14) by minimizing the combined costs to riders requesting a trip and those already traveling. The model results show that adding walking to the process reduces VHT by 10%, while riders in the highest demand areas also have to walk the farthest. The authors conclude by stating that further research is needed, such as determining the optimal vehicle fleet, how it is affected by PUDO points, and how demand further responds to these two variables.

Dynamic Ridesharing

Algorithms for approximately solving DRS problems have been extensively studied because of the difficulty of obtaining an optimal solution. In most cases in transportation literature, particularly in agent-based simulation research, this challenge leads authors to pursue simplistic heuristics for matching vehicles with trips, often placing the traveler with the closest available vehicle. It is easy to imagine scenarios where a DRS opportunity barely fulfills those constraints but is selected over another opportunity that might be a better match for the current trip, because the opportunities are not being ranked or ordered.

Alonso-Mora et al. (14) has an integer LP branch and bound solution hot started with a greedy heuristic that is fast enough to dispatch vehicles in real time. In this example, only vehicles deviated from their current paths to meet users and walking was not considered. The algorithm was formulated to minimize the wait time for passengers to be picked up. It also included a penalty for requests that went unassigned, that is, no vehicle picked up the individual who made the request. Every 30 seconds a new set of requests are analyzed, and a graph of feasible paths and vehicles that could serve them is created. If a request was not assigned during a given batch of trip matching, the penalty for not serving that request would increase. Using a fleet of 2,000 shared vehicles with a capacity of 10, wait time decreased more than in-vehicle travel time (IVTT) when compared with the base scenario of current single-occupant taxi service. 90% of rides were shared and the algorithm was light enough to analyze trips in real time, finishing assignment of each batch in less than the 30 second spacing between them. However, they also fail to adequately benchmark the algorithm against known methodologies. This makes it impossible to fully understand the model's performance and determine the validity of the model results.

Farhan and Chen (15) concluded 13 privately owned vehicles can be replaced with a single SAEV using their Capacitated Vehicle Routing Problem formulation. This solution however requires advance knowledge of the trip schedule, which in a practical context means a reservation-based system. Liu describes a quadratic formulation of the PDP that can solve small subsets of the problem and may be viable in real time adjacent to a traffic simulation (16). This approach is unique in that it solves subsets of the problem to optimality. This quadratic formulation is unique and probably the fastest solution to MVPDP available. The challenge in implementing this in an agent-based transportation simulation is generating problem subsets that are small enough to solve in real time.

METHODOLOGY

This section presents the methodology used to design simulation experiments to evaluate the impact of PUDO locations and densities on SAV fleet operation (e.g. wait time, VMT, etc.) and demand under different fare price and fleet size.

Data Preparation

For this study, the entire population of the Austin 6-county metropolitan area was simulated. This was done using the POLARIS travel demand modeling software. POLARIS is an agent-based model developed by Argonne National labs, which can model the operation of SAVs in a region (17). Similar to MATSim and other agent-based models, POLARIS enables its users to track individual vehicles through a region and its road links to individual destinations instead of aggregate zone-to-zone estimates of travel produced by trip-based models. In the model, there were 16,059 road links, 10,435 nodes, and 39,638 possible destinations created using the Capital Area Metropolitan Planning Organization's (CAMPO) 2015 roadway network. The synthetic population of 1,885,993 persons was generated using the US Census Bureau's 2018 American Community Survey (ACS) Public Use Microdata Sample (PUMS) estimates. The mode choice model was calibrated from the 2016-2017 household travel survey.

PUDO Modeling in POLARIS

This section presents the steps taken to identify and model PUDO locations in POLARIS, namely parking supply estimation, parking demand estimation and PUDO sitting, DRS, and PUDO aggregation.

Parking Supply Estimation

For PUDOs to be modeled, it was necessary to gain an understanding of where demand for curb space most challenged supply in the Austin area. When an autonomous vehicle picks up or drops off a passenger it must pull over to the curb or into a garage. Unlike some taxis or TNCs in real-world operation, an SAV will almost certainly not make stops in the middle of the roadway, no matter how brief the boarding or alighting, so as to minimize safety risks and therefore the potential liability of its owners. To calculate where this safe space was available, parking demand to supply ratios had to be calculated. Parking supply included on-street free and paid parking, and publicly-available garage or lot parking (both privately and publicly owned). Three sources of data were used.

First, the City of Austin maintains a GIS database of on-street parking locations (18). This data was corrected, and Google Street View and on-location observations were performed to fill in the gaps in on-street parking in the central business district. An example of several downtown blocks is shown in Figure 2.

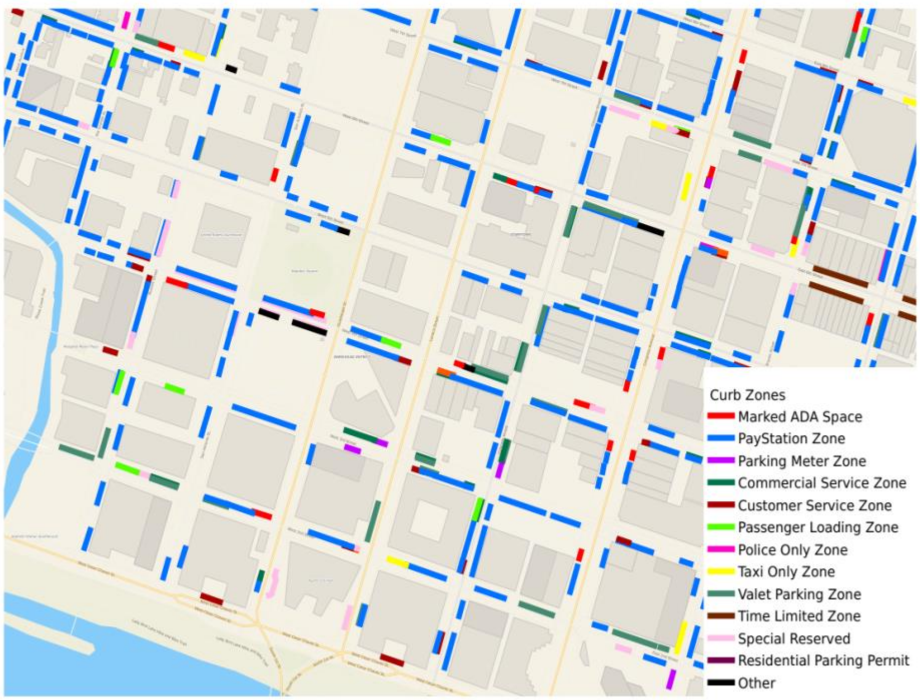

FIGURE 2 Downtown Austin curb use data

Next, an index of off-street lots and parking garages with their respective capacities geographic locations was compiled. Finally, as an exact accounting of parking on all streets in the region would be impractical and unnecessary, OpenStreetMap data was downloaded to provide a rough estimate of on-street parking in the rest of the six-county Austin metro area. Every road classified as tertiary or residential on the site was divided into 5-meter segments to conservatively calculate the supply of parking provided about every 10 meters on each side of the road on local streets.

Parking Demand Estimation and PUDO Siting

Next, demand for parking was estimated. This was done using a simulation of 25% of the population of the Austin region generated by using POLARIS. Trip ends for single occupant, carpool, and TNC trips were gathered, as these modes are the greatest competitors for pick-up and drop-off space against SAVs. Based on zones with the greatest observed demand, a geofence was created to focus SAV operation in the busiest areas, as shown in Figure 3.

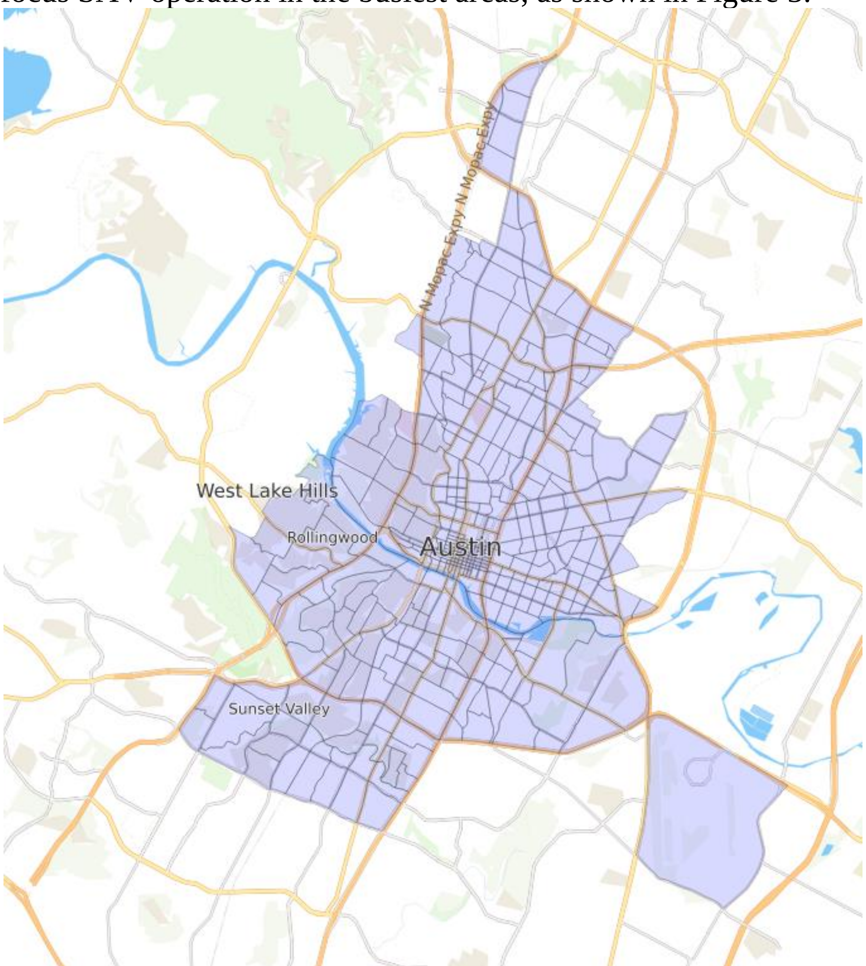

FIGURE 3 SAV service area

The total population within these traffic analysis zones (TAZs) is 326,597, while total employment is 320,262. A grid was overlaid on the data to show the relationship of trip ends to parking supply, which is a proxy for the number of trips per parking space per 24-hour weekday. The result is displayed below in Figure 4.

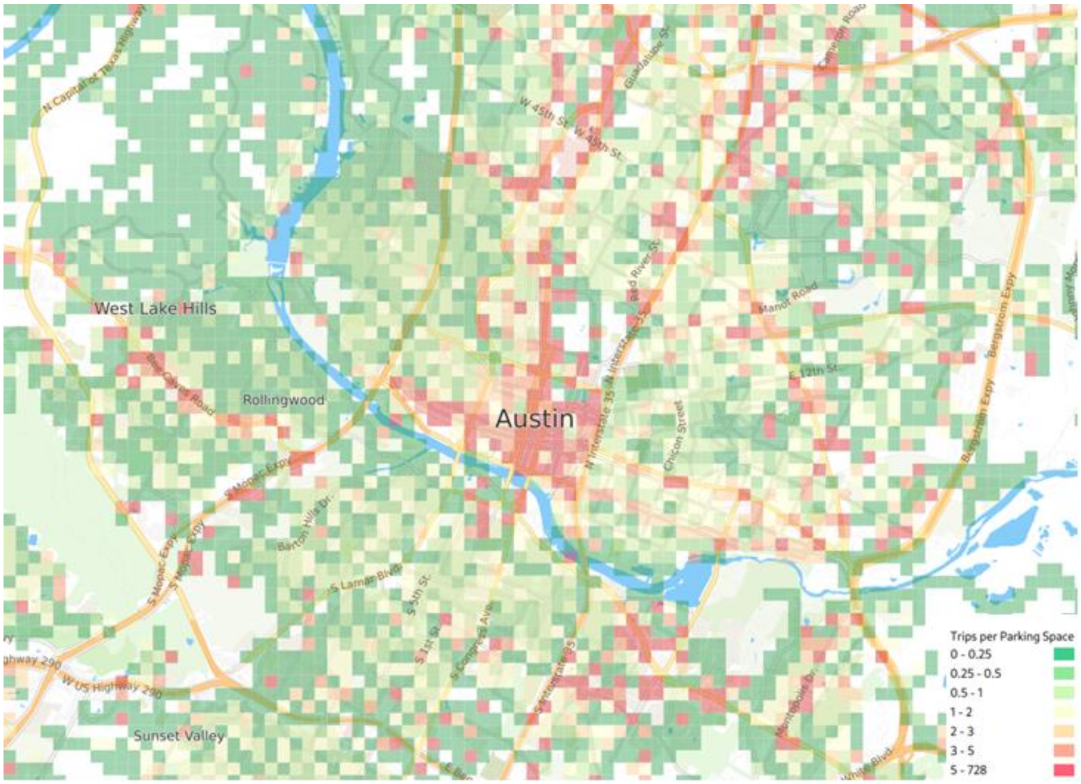

FIGURE 4 Weekday trips per parking space near the Austin, Texas CBD

Areas with greater than 5 trips per space per day in the 25% simulation, or 20 trips in a theoretical 100% simulation, are shown in red above and are a focus for this study. With such a high number of trips per space it is likely that supply will exceed demand along the curbs in that area, and SAVs will need designated PUDOs in order to find a predictably empty space to pull over in. Some of the areas in red extending from the center of the Figure 4, such as along the Lamar, I-35, Riverside, and Bee Cave corridors are a false positive for high parking utilization. This is because counts of free off-street parking provided by businesses or apartment complexes were not performed. Compared to the few dozen lots and garages in the Austin CBD, hundreds of small lots would need to be cataloged to obtain an accurate measure of parking supply along these and other similar corridors. In the future, the three sources of parking mentioned earlier could be augmented through OpenStreetMap parking catalogs or using satellite data (19). For the purposes of this simulation and as confirmed through observation of these parking lots, there is nearly always ample parking in these areas and the need for PUDOs on these stretches of road can be disregarded. Therefore, PUDO locations will be focused in the Austin CBD, traditionally defined by organizations such as the Downtown Austin Alliance as being bounded by Lamar Boulevard, Martin Luther King, Jr. Boulevard, I-35, and Town Lake (20).

DRS and PUDO Aggregation

POLARIS has been developed to include dynamic ridesharing (DRS) and PUDO aggregation. The first feature involves searching among currently occupied SAVs when a trip request is first made to the SAV operator. If a vehicle is traveling nearby and making a small diversion to pick up additional riders would not lead to significant delays for existing riders (a maximum of 5 minutes for the simulations performed in this analysis), the vehicle will be rerouted to pick up the new request. If such a vehicle is not found, the user will instead be matched to a nearby empty vehicle (21). Unlike in branded commercial implementations such as UberPool, in this simulation there is no discount for sharing a ride with others as it was assumed that sharing a vehicle would be an inherent feature of the service with no option to be the guaranteed lone occupant of a vehicle.

PUDO aggregation builds on this by directing riders to walk to designated locations to be picked up and dropped off instead of waiting at or alighting at the curb directly in front of their origin or destination. The purpose of this is threefold. First, it helps avoid the previously mentioned safety concerns and lack of curb space. Second, if stops are spaced sufficiently apart or there is enough demand in a given area, there could be enough riders boarding and getting off at a single stop to gain some of the functionality of a traditional transit system where riders are grouped together for increased efficiency. Lastly, having users walk to a convenient spot to meet their SAV can avoid the need for the vehicle to pull off the most efficient route for existing riders, such as leaving a major thoroughfare and diverting to a side street where a new rider initiated the pick-up process, and then rejoin that route again. This could lead to time savings and perhaps even pay safety dividends through reducing turn movements, though this second aspect is outside the scope of this analysis.

Model Scenarios

Three variables were adjusted to evaluate their impact on the performance of a hypothetical SAV network operating in the Austin region. The combinations of these variables produced 18 scenarios to be tested. These are shown in Table 1 below.

TABLE 1 Summary of Scenario Parameters

Parameter	Values Tested
SAV Fare	$0.50/mi
	$1/mi
SAV Fleet Size	1000 SAVs
	2000 SAVs
	4000 SAVs
PUDO Configuration	No PUDOs
	PUDOs every block in CBD
	PUDOs every third block in CBD

PUDO Spacing

Whenever a trip began or ended within ¼ mile of a PUDO (about 3 – ½ downtown Austin blocks) the rider would be instructed to walk to one of the sites to meet their assigned SAV, otherwise the SAV would pick up or drop off the passenger exactly at their origin or destination. Therefore, SAV trips could involve either door-to-door service, walking to a PUDO at the start of a trip only, being dropped off at one and walking to a destination, or using PUDOs at both ends of a trip. PUDO locations were varied between three options: no PUDO aggregation at all, one stop about every block in the CBD, or one stop about every three blocks, as shown below in Figure 3:

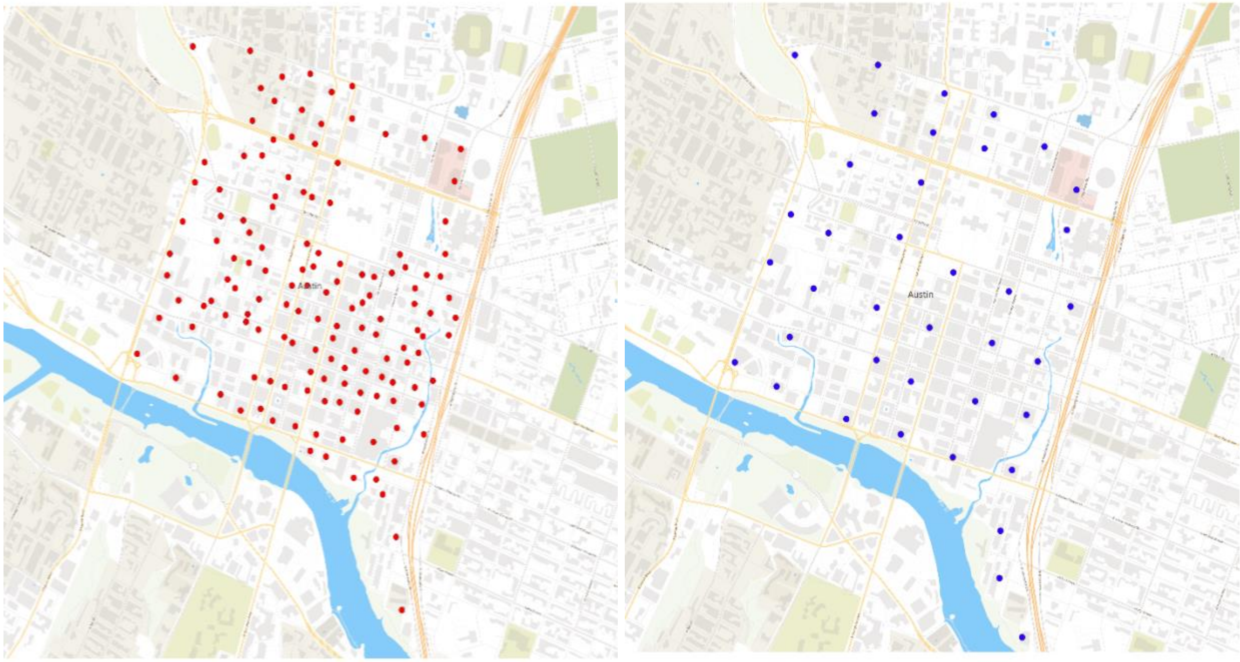

(a) (b)

FIGURE 5 (a) One-block vs. (b) two-block PUDO spacing in CBD

Currently, POLARIS does not model curb occupancy. Therefore, in the scenarios with no PUDOs the SAVs were able to instantly find a place to pull over and make their pick-up or drop-off. In the real world, SAVs must compete with other vehicles for curb areas designated as paid parking, commercial drop-off, and other designations open to SAV use. It can be assumed that the no-PUDO scenario represents a variable parking pricing policy that aims to always leave at least one space vacant on each side of a city block.

Fleet Size

The number of vehicles in the SAV fleet varied between 1000, 2000, and 4000. Each SAV had a capacity of 4 passengers. It was hypothesized that doubling or quadrupling the SAVs would not lead to an exactly corresponding increase in riders and therefore would likely reduce the operator's profits, but the extent to which ridership changed and if the number of vehicles was ever a limiting factor was a topic of interest.

Pricing

Finally, three pricing schemes for the service were implemented. First, a 50¢/mile fee, which is competitive or even slightly lower than the average cost per mile to operate a personal vehicle in the United States (22). The purpose of this was to see how many individuals could be enticed to use an SAV instead of their own car if the cost was almost the same. Second, a $1/mile fee, which was hypothesized to reduce ridership but not enough to reduce revenue per SAV.

RESULTS

All combinations of fleet sizes, fare prices, and PUDO densities were run, for a total of 18 scenarios. The most important parameters from the model results are shown below in Table 2 and Table 3.

Key Metrics

TABLE 2 Key SAV Fleet Performance Metrics, $1 Fare

Scenario	No PUDOs			1-block PUDOs			3-block PUDOs		
	1000 SAVs	2000 SAVs	4000 SAVs	1000 SAVs	2000 SAVs	4000 SAVs	1000 SAVs	2000 SAVs	4000 SAVs
SAV Trips	94,723	157,441	208,283	102,719	114,092	151,590	103,020	117,389	141,384
VMT/SAV	378	180	141	243	143	86	245	134	84
% Empty VMT	25.4	24.5	24.3	25.2	24.4	24.0	25.4	23.5	23.6
Daily Trips per SAV	94.7	78.7	52.1	107.7	57.0	37.9	103.0	58.7	35.3
Daily Revenue per SAV ($)	389.3	323.5	214.0	442.7	234.5	155.8	423.4	241.2	145.3
SAV Mode Share within Geofence (%)	10.35	17.21	22.90	10.45	12.29	14.37	9.59	12.41	13.53

TABLE 3 Key SAV Fleet Performance Metrics, 50¢ Fare

Scenario	No PUDOs			1-block PUDOs			3-block PUDOs		
	1000 SAVs	2000 SAVs	4000 SAVs	1000 SAVs	2000 SAVs	4000 SAVs	1000 SAVs	2000 SAVs	4000 SAVs
SAV Trips	100,782	132,434	248,929	109,475	132,434	152,147	118,805	121,209	143,443
VMT/SAV	364	284	184	271	151	87	257	140	82
% Empty VMT	22.6	19.1	17.5	26.5	25.9	24.5	26.0	25.2	25.9
Average Vehicle Occupancy	1.91	1.95	1.89	1.96	2.20	2.07	2.05	2.00	2.08
Daily Trips per SAV	100.8	66.2	62.2	102.7	57.0	32.4	102.2	57.9	31.4
Daily Revenue per SAV ($)	207.1	136.1	127.9	211.1	117.2	66.6	210.0	119.1	64.6
SAV Mode Share within Geofence (%)	10.91	17.86	26.56	11.09	12.41	14.30	11.01	12.60	13.84

Fleet Size

As expected, increasing the fleet size always leads to increased ridership no matter the combination of the other two variables. Figure 6 shows the trips per hour for an SAV network with no PUDOs, a 50¢ fare, and different SAV fleet sizes (1000, 2000 and 4000), which showed the greatest variation of ridership among the scenarios.

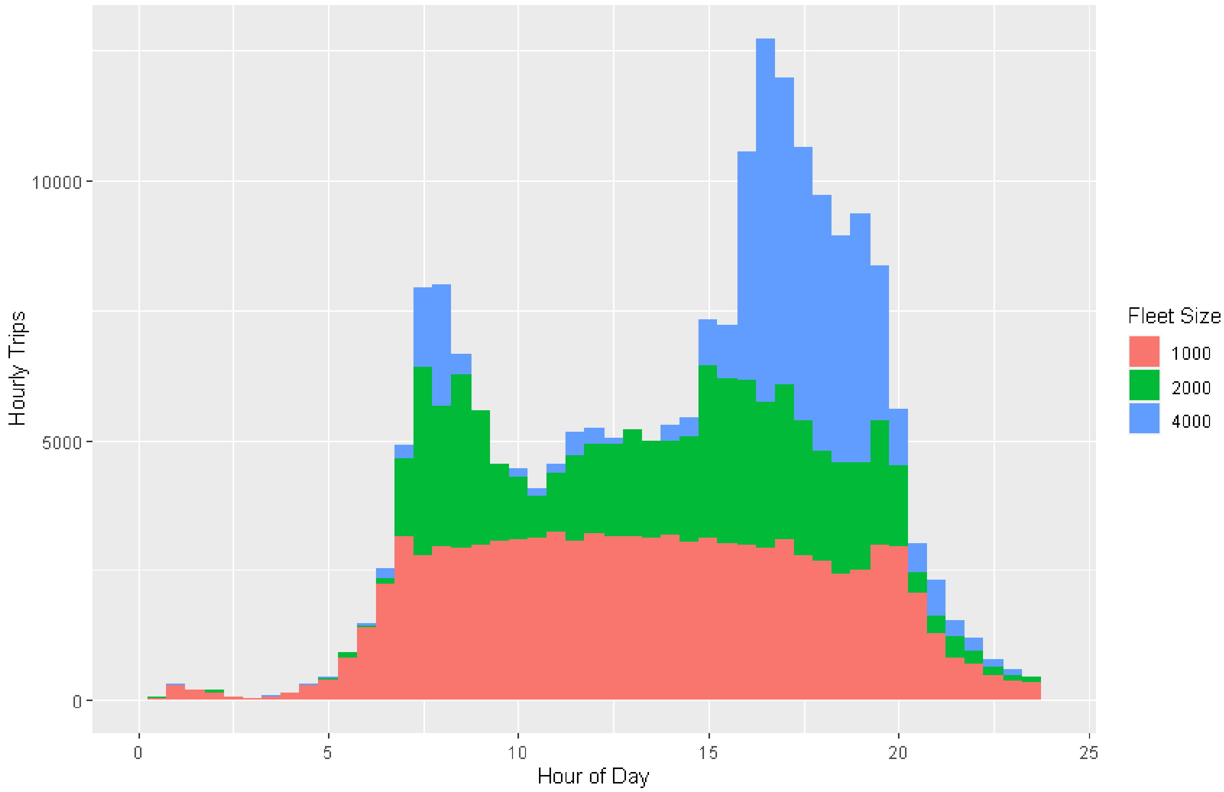

FIGURE 6 SAV weekday trips per hour for varying fleet sizes, half-hour bins

This figure shows that both the 2000-vehicle and 4000-vehicle fleets manage to accommodate demand during the middle of the day, while the 1000-vehicle fleet peaks at about 2500 trips per half hour. The 4000-vehicle fleet is especially able to accommodate additional demand as the PM peak goes on while the smaller fleets are unable to keep up, possibly through a physical lack of space inside the vehicles to add more passengers, long wait times discouraging others from choosing the SAV mode, or other means. The exact operating statistics of the SAVs and the Austin city core, such as vehicle occupancy at each hour hour and delay on links in the CBD, are interesting factors that should be explored in the future to explain the exact causes of the lower ridership.

Vehicle-miles traveled (VMT) and empty vehicle-miles traveled (eVMT) did slightly decrease as fleet size increased. As more vehicle saturated the served region, they were therefore on average closer to new trip requests and had to drive a shorter distance with no passengers.

Interestingly, there was no general pattern for average vehicle occupancy, which was maintained at about 2 passengers per SAV during revenue trips. If occupancy could be increased from its present equilibrium and approach the 4-passenger capacity, this could theoretically lead to increased ridership. However, additional dynamic ridesharing would lead to longer trip times for users and the SAV mode would therefore become less attractive.

Pricing

The last three parameters are especially relevant for the commercial or financial feasibility of operating an SAV fleet. Fleet operators must ensure that their investment in SAVs is being put to good use through a high average occupancy, as vehicles sitting unused mean money going to waste. The daily trips per SAV begins with quite high values, about 100 per day, then decreases to nearly 30 trips per day as fleet size grows. SAV mode share could certainly increase beyond 25%, but it might not necessarily be cost effective to operate a larger fleet. The point where marginal revenue and costs per SAV cross over is an important factor in determining how many trips will be made by SAVs. Two competing goals exist for an SAV network, maximizing ridership and maximizing profits, and at market equilibrium the total welfare for users and riders is maximized. A municipality could perhaps subsidize a private SAV network or operate its own fleet in order to increase total consumer welfare, serving a greater number of users. Fare prices did have a slight effect on ridership, with more riders for the lower-cost option, but per-vehicle revenue was higher for the more expensive SAV pricing.

PUDO Density

As PUDOs were implemented, there was a sharp drop-off in ridership with 1-block PUDOs, then a more minor drop when PUDOs moved to being located every three blocks. This shows that in an idealized world where SAVs could always make door-to-door trips, they would achieve a greater mode share and reduce private vehicle traffic. However as curb space, or the lack thereof, must be taken into account in operating SAVs, the peak demand on curbs by SAVs must be measured.

First, an assumption must be made on the maximum capacity of a PUDO per hour. Assuming that an SAV takes 30 seconds to pick up or drop off a passenger, the theoretical hourly capacity of a PUDO would be 120 SAVs. Hence a PUDO location with 135 trip ends during its busiest hour would require two spaces to handle peak trips. This ignores the fact that trips are not perfectly spaced throughout the hour, but will suffice for the illustrative purposes of the next table. In Table 4, the distribution of required SAV spaces at each PUDO site is shown for the 1- and 3-block versions of the $1/mile, 4000-vehicle fleet.

TABLE 4 Distribution of Required Curb Spaces at Each PUDO Site

Number of Spaces	1-block PUDO Spacing	3-block PUDO Spacing
1	64%	25%
2	17.0%	27.5%
3	8.1%	15%
4	4.4%	7.5%
5+	5.9%	25%
Total Spaces	263	160

As shown in Table 4, over 1/3 of PUDO locations would need more than one curb space when there is a PUDO on every block, and this jumps to 75% when PUDOs are spaced out to every three blocks. The estimate also shows one of the advantages of increased spacing between PUDOs; the total number of spaces required to accommodate peak traffic is reduced by 39%. There is always a tradeoff to be had between SAV ridership and the amount of curb spaces available for other uses in the CBD. If there is only a PUDO location every 3 blocks, curb space is freed up on other blocks for paid parking, transit, bicycle parking, or other uses. However, the increased walk access time leads to reduced ridership. The table also shows that since 94% of blocks need four or less PUDO spaces in the 1-block scenario, one space could be placed on each side of the block to achieve service similar to the scenario with no PUDO aggregation.

CONCLUSIONS

This study sought to analyze the effect of various configurations of fleet size, pricing, and PUDO spacing on the performance of SAVs. It was shown that greater fleet sizes accommodated more passengers, especially at peak hours, and decreased eVMT, but led to reduced revenue per vehicle. Moving from no PUDOs to PUDOs on every block, and then to a less-dense configuration, caused SAV mode share to decrease. However, the total number of curb spaces required for PUDOs decreased by almost 40%. Greater per-mile pricing also reduced ridership, but per-vehicle revenue increased.

This analysis had some limitations, such as imperfect mode choice and pedestrian walking distance models. The current mode logit model suffered from a small number of taxi/TNC trips in the travel survey used to estimate it, leading to a large alternative-specific constant for the SAV mode and few statistically-significant additional parameters. Because of this, it is difficult to determine a realistic mode share and SAV demand on the network, and some of the results of the simulation may have been due to random noise. In addition, empty SAVs currently wait at their last trip end, without taking up any space on the curb, until they are next assigned to a trip. A better simulation of the real world not allow SAVs to stop at a PUDO occupied by an SAV that is empty or performing pick-up/drop-off, and/or force all empty SAVs to wait outside the CBD on local streets with low curb occupancy. Overall, it is clear that POLARIS provides a powerful modeling framework with a plethora of parameters that can be adjusted to view their effect on travel behavior. Future work will involve leveraging this framework to mitigate the limitations expressed above.

In conclusion, as cities evolve and integrate new means of transportation, such as micromobility and SAVs, their curb usage must evolve as well. To safely accommodate these new modes, some space must be reclaimed from existing uses such as paid on-street parking. When planning an SAV network, the density of PUDO locations has an enormous impact on mode share as well as the amount of curb real estate required for them. Not using dedicated PUDO locations at all might theoretically lead to better SAV mode share, but realistically at least some PUDOs must be implemented. Once PUDO sites must be chosen, their spacing is a crucial decision that affects both SAV ridership and the amount of curb space that must be taken away from other uses.

ACKNOWLEDGEMENTS

This research is supported by the Ford Motor Company under the University Research Program (URP). The authors would like to thank Richard Twumasi-Boakye, Archak Mittal, and Andrea Broaddus, R&A, Ford Motor Company for their guidance and support in this project.
The authors would like to give special thanks to Matthew Dean for assistance with the literature review of this paper, as well as Dorcas Olayoe for compiling the database of off-street garages and parking lots in downtown Austin. In addition, we would like to thank the team at Argonne National Laboratory for continually developing POLARIS, especially Krishna Murthy Gurumurthy for his assistance with coding. Finally, we thank the Texas Advanced Computing Center for allowing us to run the simulations on their machines.

REFERENCES

1. Marsden, Greg, and Iain Docherty. "Parking futures: Curbside management in the era of 'new mobility' services in British and Australian cities." Land Use Policy 91 (February 2020): 104012. https://doi.org/10.1016/j.landusepol.2019.05.031.

2. Butrina, P., Le Vine, S., Henao, A., Sperling, J., Young, S.E., 2020. Municipal adaptation to changing curbside demands: Exploratory findings from semi-structured interviews with ten U.S. cities. Transp. Policy 92, 1–7. https://doi.org/10.1016/j.tranpol.2020.03.005.

3. Małecki, Krzysztof. "A computer simulation of traffic flow with on-street parking and drivers' behaviour based on cellular automata and a multi-agent system." Journal of Computational Science 28 (September 2018): 32-42. https://doi.org/10.1016/j.jocs.2018.07.005.

4. Levin, Michael W., Kara M. Kockelman, Stephen D. Boyles, and Tianxin Li. "A General Framework for Modeling Shared Autonomous Vehicles with Dynamic Network-Loading and Dynamic Ride-Sharing Application." Computers, Environment and Urban Systems 64 (July 1, 2017): 373–83. https://doi.org/10.1016/j.compenvurbsys.2017.04.006.

5. Bösch, Patrick M., Francesco Ciari, and Kay W. Axhausen. "Autonomous Vehicle Fleet Sizes Required to Serve Different Levels of Demand." Transportation Research Record 2542, no. 1 (January 1, 2016): 111–19. https://doi.org/10.3141/2542-13.

6. Fagnant, Daniel J., Kara M. Kockelman, and Prateek Bansal. "Operations of Shared Autonomous Vehicle Fleet for Austin, Texas, Market." Transportation Research Record: Journal of the Transportation Research Board 2536 (January 1, 2015): 98–106. https://doi.org/10.3141/2536-12.

7. Hörl, Sebastian. "Agent-Based Simulation of Autonomous Taxi Services with Dynamic Demand Responses." Procedia Computer Science 109 (January 1, 2017): 899–904. https://doi.org/10.1016/j.procs.2017.05.418.

8. Bischoff, Joschka, and Michal Maciejewski. "Simulation of City-Wide Replacement of Private Cars with Autonomous Taxis in Berlin." Procedia Computer Science 83 (January 1, 2016): 237–44. https://doi.org/10.1016/j.procs.2016.04.121.

9. Loeb, Benjamin, and Kara M. Kockelman. "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 (March 1, 2019): 374–85.

10. Kockelman, Kara, Stephen Boyles, Christian Claudel, Peter Stone, Lisa Loftus-Otway, Purser Sturgeon, Guni Sharon, et al. "Phase 2 - Bringing Smart Transport to Texans: Ensuring the Benefits of a Connected and Autonomous Transport System in Texas." TxDOT, July 2018. https://trid.trb.org/view/1526257.

11. Stead, Dominic, and Bhavana Vaddadi. "Automated vehicles and how they may affect urban form: A review of recent scenario studies." Cities 92 (September 2019): 125-133. https://doi.org/10.1016/j.cities.2019.03.020.

12. Li, Rong-Hua, Lu Qin, Jeffrey Xu Yu, and Rui Mao. "Optimal Multi-Meeting-Point Route Search." IEEE Transactions on Knowledge and Data Engineering 28 (March 1, 2016): 770-784. https://doi.org/10.1109/tkde.2015.2492554.

13. Fielbaum, Andres, Xiaoshan Bai, and Javier Alonso-Mora. "On-demand ridesharing with optimized pick-up and drop-off walking locations." Transportation Research Part C: Emerging Technologies 93, 2021. https://doi.org/10.1016/j.trc.2021.103061.

14. Alonso-Mora, Javier, Samitha Samaranayake, Alex Wallar, Emilio Frazzoli, and Daniela Rus. "On-Demand High-Capacity Ride-Sharing via Dynamic Trip-Vehicle Assignment." Proceedings of the National Academy of Sciences 114, no. 3 (January 17, 2017): 462–67. https://doi.org/10.1073/pnas.1611675114.

15. Farhan, J., and T. Donna Chen. "Impact of Ridesharing on Operational Efficiency of Shared Autonomous Electric Vehicle Fleet." Transportation Research Part C: Emerging Technologies 93 (August 1, 2018): 310–21. https://doi.org/10.1016/j.trc.2018.04.022.

16. Liu, Zhiguang, Tomio Miwa, Weiliang Zeng, Michael G. H. Bell, and Takayuki Morikawa. "Dynamic Shared Autonomous Taxi System Considering On-Time Arrival Reliability." Transportation Research Part C: Emerging Technologies 103 (June 1, 2019): 281–97. https://doi.org/10.1016/j.trc.2019.04.018.

17. Gurumurty, Krishna Murthy and Kara Kockelman. "How Much Does Greater Trip Demand and Aggregation at Stops Improve Dynamic Ride-Sharing in Shared Autonomous Vehicle Systems?" Bridging Transportation Researchers Conference (August 5, 2020). https://www.caee.utexas.edu/prof/kockelman/public_html/TRB21savstopaggregationbloomington.pdf.

18. Austin Transportation Department – Parking Enterprise Division. "Parking Inventory" [Map]. (January 12, 2021). https://services.arcgis.com/0L95CJ0VTaxqcmED/ArcGIS/rest/services/Parking_Inventory/FeatureServer.

19. Hawkins, Jason. "Modelling Spatial Location Choice and Transition for a Changing Urban Landscape." University of Toronto (June 2021). https://hdl.handle.net/1807/106371.

20. Downtown Austin Alliance, "About Us". (2021) https://downtownaustin.com/about-us/.

21. Gurumurthy, Krishna Murthy, Felipe de Souza, Annesha Enam, and Joshua Auld. "Integrating Supply and Demand Perspectives for a Large-Scale Simulation of Shared Autonomous Vehicles." Transportation Research Record: Journal of the Transportation Research Board 2674 (June 2, 2020): 181-192. https://doi.org/10.1177/0361198120921157.

Bureau of Transportation Statistics. "Average Cost of Owning and Operating an Automobile." (2019). https://www.bts.gov/content/average-cost-owning-and-operating-automobilea-assuming-15000-vehicle-miles-year.