Autonomous or Driver-less Vehicles: Implementation Strategies and Operational Concerns

Transportation Research Part E

Neda Masoud
Assistant Professor
Department of Civil and Environmental Engineering
University of Michigan, Ann Arbor 48109
nmasoud@umich.edu
Corresponding author

R. Jayakrishnan
Professor
Department of Civil and Environmental Engineering
University of California, Irvine Irvine, CA, USA, 92697
rjayakri@uci.edu

Abstract

Autonomous vehicles are expected to shift not only the driving paradigms but also the notion of vehicle ownership. Although autonomous vehicles are believed to introduce many safety, mobility, and environmental benefits, they will be initially priced relatively highly. This paper assesses the potential for circumventing this barrier by promoting a shared ownership program in which households form clusters that share the ownership and ridership of a set of autonomous vehicles. Such a program will increase the utilization rate of vehicles, making ownership of autonomous vehicles more economical. We study parameters that affect the benefits expected from autonomous vehicles, and introduce policy directions that can boost these benefits.

1 Introduction

Sharing economy, also known as collaborative consumption, is a fairly old concept that focuses on the benefits obtained from sharing resources (products or services) that would otherwise go unused. Although communities have been using the concept of sharing economy locally for many years, advent of internet has led to its spread in global markets, and has highlighted its benefits.

The sharing economy model has been historically used for high-value commodities, such as exotic automobiles, yachts, private jets, vacations homes, and the like. Although it has been long realized that taking ownership of under-utilized high-value assets may not be always economically viable, this economic model has become more popular recently for less expensive resources as well, thanks to new platforms that allow easy and quick development of companion mobile applications. Sharing economy in the context of mobility should not be confused with Mobility as a Service (MaaS). While using MaaS is typically attributed to foregoing vehicle ownership altogether and outsourcing transportation as a service, shared-use mobility does not necessitate foregoing vehicle ownership and can take various forms (e.g., public transportation, ridesharing, carsharing, bikesharing), where some players may own vehicles and share vehicles/rides with others.

Autonomous (also known as driver-less and self-driving) vehicles are expected to enter the market in the near future. Although these vehicles introduce many benefits such as higher degrees of safety and mobility to the users, their high prices, primarily resultant from the added cost of different types of sensors they need to be equipped with, can be prohibitive when it comes to purchasing them (Wickerham, 2017). On the other hand, autonomous vehicles can reduce the number of vehicles a household may need to perform daily tasks, since these vehicles can drive themselves to locations where there is demand for transportation. One possible strategy to make autonomous vehicles more affordable is to encourage shared ownership of these vehicles. Moreover, it is possible to lower the ownership cost of autonomous vehicles further by (*i*) promoting shared ridership of these vehicles, and (*ii*), renting out these vehicles when they are not being used by their owners.

The ability of autonomous vehicles to operate without a driver lowers the number of vehicles a household may require to perform its daily tasks. Figure 1 demonstrates the average daily vehicle miles traveled (VMT) by each vehicle in a household in the US in 2009. This figure suggests that the higher the number of vehicles owned by a household, the less the average utilization rate of each additional vehicle tends to be. Although for a typical household owning more than one vehicle might be financially justifiable considering the level of comfort and peace of mind vehicle ownership may bring, this justifiablity decreases with the purchase of additional vehicles. Apart from the initial investment (or monthly payments), the cost of insurance, depreciation of value, and parking can turn vehicle ownership into a financial burden. With autonomous vehicles, fewer vehicles can cover the same set of trips compared to a higher number of legacy vehicles, making the idea of each household member owning a car an obsolete one.

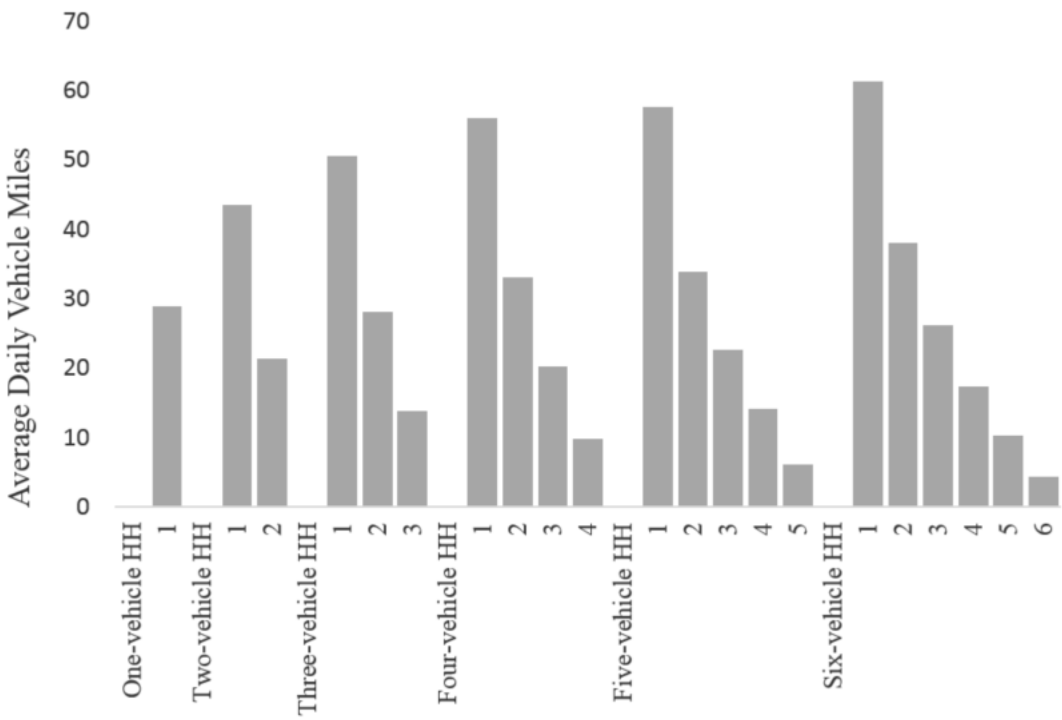

Figure 1: Average daily vehicle miles for households with various number of vehicles. (Data obtained from the National Household Travel Survey (NHTS), 2009.)

Though no rigorous analysis under optimal operations appear to have been done in the literature yet, there is some awareness of such possibilities in the automotive industry as well as among the increasing number of researchers and aficionados of autonomous vehicles (Schall, 2015; Naughton, 2015; Fagnant and Kockelman, 2014; Schoettle and Sivak, 2015). In this paper, we introduce the shared vehicle ownership and ridership (SVOR) program in which a group of households jointly own and use a set of autonomous vehicles. Households can share rides with each other if the spatiotemporal distributions of their trips allow for it. We propose analytical optimization schemes to study the impact of this program at both individual and system levels.

2 Literature Review

The proposed program in this paper combines three shared mobility models: fractional vehicle ownership, and peer-to-peer car- and ride-sharing. We combine the individual advantages offered by each of these models, and propose a system that aims at maximizing efficiency.

Fractional ownership of luxury commodities emerged in the US in 1970's with real estate time shares, and was later spread to other high value commodities (Garigliano, 2007). Ford, in partnership with Zoomcar, was the first company to start a pilot project of fractional ownership of non-luxury vehicles in Bengaluru, India, as a part of its 25 mobility experiment initiative (John and Phadnis, 2015). During this three month pilot which took place in 2015, Zoomcar provided 5 vehicles, each of which were shared by 6 individuals. This was a first step for Ford and Zoomcar to study the impacts and implications of fractional vehicle ownership. Follwoing this experiment, Ford made a $24M investment in Zoomcar in 2016 (Prasad, 2016).

In the proposed system in this paper, households who share the ownership of vehicles have the possibility of sharing rides, if their trips are spatiotemporally compatible. Ridesharing systems are well-studied in the literature, and a large volume of studies have confirmed their numerous benefits, including reduction in vehicle miles traveled (VMT), and less damage to the transportation infrastructure and the environment (Chan and Shaheen, 2012; Heinrich, 2010).

Literature suggests that when it comes to switching to autonomous vehicles, previous modality preferences play an important role. Individuals whose previous modality relies heavily on private vehicles are less likely to forego vehicle ownership in favor of a central shared mobility service provider (Krueger et al., 2016). Shared ownership of autonomous vehicles can provide such households with the peace of mind of vehicle ownership, while at the same time leaving the door open for shared ridership.

In order to make shared ownership of autonomous vehicles more affordable, households can rent their vehicles when they are not being used. The advantage of using autonomous vehicles in carsharing programs is that the complicated dispatching problem that one-way car sharing systems face does not arise, since autonomous vehicles can drive themselves and there is no need to plan and dispatch drivers for their re-distribution in the network. A comprehensive review and classification of carsharing models can be found in Barth and Shaheen (2002). There are multiple studies that link carsharing to reduction in household vehicle holdings, increase in older vehicle sales, and postponing vehicle purchase (Martin et al., 2010; Walb et al., 1986). Additional studies highlight the positive impact of carsharing on VMT (Cervero and Tsai, 2004; Cervero et al., 2007). There are studies in the literature that have looked at on-demand carsharing systems with autonomous vehicles. Fagnant et al. (015b), for example, use simulations to study the implications of short-term on-demand autonomous vehicle rentals in the Austin, Texas area. Fagnant and Kockelman (015a) look at the problem of finding the optimal fleet size for shared autonomous vehicle deployment in Austin, Texas.

To the best of our knowledge, this study is the first to focus on shared ownership and ridership of autonomous vehicles with an analytical formulation. In order to implement the SVOR program, we need to form clusters of households, where members of each cluster jointly own a set of vehicles. The goal is to increase efficiency by finding the minimum number of vehicles each cluster requires, and allowing members of each cluster to rideshare if the opportunity presents itself. This is similar to the problem of finding the minimum number of vehicles in a dial-a-ride problem (DARP) with time windows (Jaw et al., 1986; Psaraftis, 1983). In addition, we allow clusters of households to rent out their vehicles using a central carsharing system. The problem of allocating vehicles to dynamic requests bears similarities to the dynamic DARP with time windows. In this paper, we formulate this problem as a pure binary program in a time-expanded network, and demonstrate through an experimental study that this binary formulation can efficiently solve the targeted problem.

Cordeau and Laporte (2007) provide a literature review on the algorithms developed for the dynamic DARP. The computational time, and/or number of requests these algorithms are able to manage poses significant limitations on the system proposed in this paper. We propose a greedy heuristic algorithm that is able to provide high quality solutions within a short period of time. Finally, we implement the SVOR program for a sample of households in San Diego, California, and comment on the resulting efficiency at the individual and system levels.

The contributions of this paper are three-fold. First, we introduce a general mathematical framework to model a system with various levels of shared autonomous vehicle ownership and ridership. It is believed that upon deployment, autonomous vehicles can enhance safety and mobility and curb environment side-effects of the transportation sector. The second contribution of this study is to put a question mark in front of this widely accepted premise, and point out factors that could play a role in the degree to which, and the circumstances under which, these potential benefits can be realized. Finally, we demonstrate the degree to which a shared ownership and ridership model can reduce the number of vehicles needed to satisfy the transportation needs of a population, thereby reducing the cost of autonomous transportation through fractional vehicle ownership.

3 SVOR: Shared Vehicle Ownership and Ridership

Envision a set of households, F, who share the ownership of a set of autonomous vehicles, V. These households form a cluster to which the vehicles under their shared ownership belong. Each vehicle $v \in V$ has the capacity to carry C_v number of passengers. Each household $f \in F$ has a set of essential trips that need to be served by autonomous vehicles. We define set M_e to include all the essential trips of the households that belong to a cluster. Conceptually, essential trips are trips for which individuals need to have timely and regular access to vehicles; however, households can include any other types of trips in the set of essential trips if they need to ensure regular and timely access to a vehicle for such trips. Common types of essential trips may include work-based trips, grocery shopping, and trips to school.

For a given trip k, a cluster member needs to input into the system the location of the origin of the trip, OS_k, the location of the destination of the trip, DS_k, the earliest departure time from the origin location, ED_k, and the latest arrival time at the destination location, LA_k. While vehicles are idle, they can be rented out to satisfy a set of on-demand transportation requests, M, in order to cover a part of the system cost. A rental request $k \in M$ should include the location where a vehicle needs to deliver itself (OS_k), and the location where it needs to return (DS_k), along with the rental period duration (P_k), and the rental time window ($[ED_k, LA_k]$).

The first goal of the system is to advise households in a cluster on the optimal number of vehicles they need to purchase to cover their set of essential trips. In the interest of higher efficiency, the system is designed to allow cluster members to rideshare, if the spatiotemporal proximity of their trips permits it. The second goal of the system is to maximize the total number of on-demand car rentals in order to maximize the external revenue generated. These goals are implemented sequentially, i.e., we first determine the optimal number of vehicles for each cluster of households (under different clustering schemes), and then use these vehicles for carsharing during their idle times. In the next section, we mathematically model these two problems.

3.1 Mathematical Modeling

The SVOR program finds the min number of vehicles for clusters of households to serve their set of essential trips. When not serving the essential trips of cluster members, these vehicles can be rented out to the general public. The specifications of the number of vehicles needed by each cluster, the itineraries of these vehicles, and how these vehicles should be used to serve on-demand rental requests can be obtained by solving two mathematical problems. In this section, we elaborate on the properties of these mathematical formulations.

In order to model the system defined in the previous section, we formulate two optimization problems. The first problem presented in section 3.2 finds the optimal number of autonomous vehicles that should be owned by a cluster under the constraint that the cluster's set of essential trips be served. Furthermore, this optimization problem generates itineraries for all essential trips. The result would be a system best described as a *personalized transit system* for each individual - Pick-up and drop-off times and locations would be pre-determined and fixed, but personalized to each individual's travel needs. The second problem presented in section 3.3 uses the vehicles' idle times to serve the maximum number of on-demand car rental/trip requests.

We will formulate these two optimization problems on a time-expanded network, where a link identifies the origin and destination of a trip as well as its departure and arrival times. Formulating

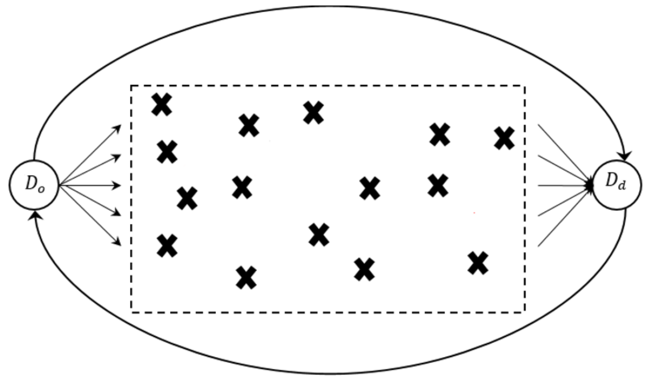

Figure 2: A typical network to demonstrate the connection of depot stations together, and to members of set S

the problem on a time-expanded network, rather than using a more general dial-a-ride formulation, enables us to exploit the unique structure of the problem resulting from the fact that trips (requested by individuals) and vehicles are both available within narrow sprites in time and space. This simple observation enables us to simply filter out the links that do not satisfy the spatio-temporal constraints imposed by trip time windows and vehicle availability constraints. This more than compensates the larger size of the link sets resulting from increasing the dimensionality of links by adding time tags. Furthermore, formulating the two optimization problems on a time-expanded network obviates the need for imposing explicit time constraints on trips by defining sets that only include links that readily satisfy these constraints. Additionally, adding a time component to the definition of links renders tour elimination constraints, which substantially increase the complexity of pick-up and delivery problems, unnecessary.

Aside from reducing the number of constraint sets, formulating the problems on a time-expanded network enables us to form pure binary (zero-one) programs. The resulting optimization problems have very tight linear relaxations, limiting the need for branching and bounding to only very specific cases, where the spatiotemporal proximity between trips and vehicles is significant.

For a given cluster, let us define the set of stations, denoted by S_e. This set contains the origin and destination locations of the cluster's essential trips. Furthermore, let us define set S to contain all the origin and destination locations of all essential and non-essential trips (by all clusters). By introducing stations, we discretize the space dimension of the problem. In addition, let us discretize the study time horizon into a set of short time intervals with length δt. We define set T to include all time intervals in the study time horizon. In this study, we use 5-min time intervals. In a network discretized in both time and space, let us define a node n as a tuple $(t_i, s_i) \in T \times S$. Consequently, we can define a link ℓ as a tuple of nodes $\ell = (n_i, n_j) = (t_i, s_i, t_j, s_j)$, where $(t_j - t_i)\delta t$ is the travel time between stations s_i and s_j. Let us define set L to include all links in the network.

We define an origin depot, D_o, from which all cluster vehicles depart in the beginning of the day, and a destination depot D_d, to which all cluster vehicles return at the end of the day. The depots are virtual stations used to assist in the formulation of the problem. The origin depot, D_o, is connected to all stations in set S, and all stations in S are connected to the destination depot, D_d. Finally, D_o and D_d are connected to each other. Figure 2 displays a typical network structure and demonstrates the connection between the depots and the stations.

We use the pre-processing procedure presented by Masoud and Jayakrishnan (2017) to generate the set of links that are reachable by each trip, given its origin and destitution as well as its time window. Let us denote by L_k the link set for trip $k \in \{M \cup M_e\}$.

3.2 Routing of Autonomous Vehicles

As stated previously, the SVOR program requires solving two mathematical problems. The first probelem pesented in this section finds the optimal number of vehicles to serve a cluster's set of essential trips. This problem is formulated in the problem set (3). The formulation requires two sets of decision variables defined in (1) and (2).

$$x_\ell^v = \begin{cases} 1 & \text{If vehicle } v \text{ travels on link } \ell \\ 0 & \text{Otherwise} \end{cases} \tag{1}$$

$$y_\ell^{kv} = \begin{cases} 1 & \text{If trip } k \text{ is carried out by vehicle } v \text{ on link } \ell \\ 0 & \text{Otherwise} \end{cases} \tag{2}$$

Assuming that the households in a cluster have a total of m members, $|m|$ serves as an upper-bound on the number of vehicles needed to serve the cluster. Therefore, to determine the minimum number of vehicles for a given cluster, we formulate an optimization problem, assuming there to be $|m|$ vehicles available, and try to maximize the number of vehicles that are not used.

Constraint sets (3b) and (3c) force all vehicles to go back from D_d to D_o at the end of the day. Constraint set (3d) is the flow conservation constraint, forcing all vehicles that enter a station at a given time interval to leave that station at the same time interval. Notice that vehicles do not have to physically leave a station; this can be represented by links that have similar origin and destination stations, but different departure and arrival time intervals. Constraint sets (3e)-(3g) route the set of trips in the network. Constraint set (3e) and (3f) ensure that a trip leaves its origin station and enters its destination station within its time window, respectively. Constraint set (3g) is the flow conservation constraint for trips. Constraint set (3h) serves two purposes: it links vehicle routes to trip routes, and ensures that the number of individuals assigned to each vehicle at any moment in time does not exceed the vehicle's capacity, where C_v is capacity of vehicle v.

Since all vehicles have to start at the origin depot and terminate at the destination depot, vehicles that are excessive and are not actually routed in the system have to take the direct link connecting D_o to D_d. Therefore, vehicles who travel on this link are not actually being used. Consequently, to minimize the number of used vehicles, we maximize the vehicles that travel on this link, as mathematically stated in the objective function of the problem in (3). The second term in the objective function minimizes the total travel time by vehicles in the network. We set a positive weight W for the first term in the objective function to take into account the relative importance of minimizing the number of vehicles in a cluster compared to the total travel time experienced by the cluster members. The solution to this problem simultaneously provides the minimum number of vehicles required to serve the essential trips and itineraries for the trips and the vehicles.

While travel times in model (3) are considered constant, it is perceivable that with a high penetration rate this program could affect network-level travel times. Under such circumstances, a bi-level model can be used, where the mathematical program presented in (3) acts as the upper-level problem. The lower-level problem takes the trip table generated by the upper-level problem as an input, and generates updated travel times (using simulations or based on a user-equilibrium framework), which can then be fed back to the mathematical model in the first level. The model iterates until no further changes in travel times can be observed. Gao et al. (2005) provide a solution methodology for solving such a bi-level model. Note that with a high penetration rate, failure to account for changes in travel times would lead to misleading results that do not hold in practice.

$$\text{Minimize} \quad -W \sum_{\substack{v\in V,\ t_i\in T,\ t_j\in T:\\ \ell=(t_i,D_o,t_j,D_d)\in L}} x_\ell^v + \sum_{\substack{v\in V,\ell\in L:\\ s_i\neq s_j}} x_\ell^v \tag{3a}$$

$$\text{Subject to:} \quad \sum_{\substack{\ell=(t_i,s_i,t_j,s_j)\in L:\\ s_i=D_d,s_j=S_e\backslash D_o}} x_\ell^v = 0 \qquad \forall v \in V \tag{3b}$$

$$\sum_{\substack{\ell=(t_i,s_i,t_j,s_j)\in L:\\ s_i=D_d,s_j=D_o}} x_\ell^v = 1 \qquad \forall v \in V \tag{3c}$$

$$\sum_{\substack{s_i\in S_e,t_i\in T:\\ \ell=(t_i,s_i,t,s)\in L}} x_\ell^v = \sum_{\substack{s_j\in S_e,t_j\in T:\\ \ell=(t,s,t_j,s_j)\in L}} x_\ell^v \qquad \forall v\in V, t\in T, s\in S_e\backslash D_o: \ \ell=(t_i,s_i,t,s)\in L \tag{3d}$$

$$\sum_{v\in V}\sum_{\substack{\ell\in L_k:\\ s_i=OS_k}} y_\ell^{kv} - \sum_{v\in V}\sum_{\substack{\ell\in L_k:\\ s_j=OS_k}} y_\ell^{kv} = 1 \qquad \forall k\in M_e \tag{3e}$$

$$\sum_{v\in V}\sum_{\substack{\ell\in L_k:\\ s_j=DS_k}} y_\ell^{kv} - \sum_{v\in V}\sum_{\substack{\ell\in L_k:\\ s_i=DS_k}} y_\ell^{kv} = 1 \qquad \forall k\in M_e \tag{3f}$$

$$\sum_{\substack{s_i\in S_e,t_i\in T:\\ \ell=(t_i,s_i,t,s)\in L_k}} y_\ell^{kv} = \sum_{\substack{s_j\in S_e,t_j\in T:\\ \ell=(t,s,t_j,s_j)\in L_k}} y_\ell^{kv} \qquad \forall v\in V, k\in M_e, t\in T, \ s\in S_e\backslash\{OS_k, DS_k\}: \ \ell=(t_i,s_i,t,s)\in L_k \tag{3g}$$

$$\sum_{k\in M_e,\ell\in L_k} x_\ell^{kv} \leq C_v x_\ell^v \qquad \forall v\in V, \ell\in L \tag{3h}$$

$$x_\ell^v, y_\ell^{kv} \in \{0,1\} \tag{3i}$$

3.3 On-demand Carsharing

The pick-up and drop-off schedules for the set of essential trips are determined by the optimization problem in the previous section. The idle time windows of vehicles can be used to serve on-demand transportation requests through a central car rental service provider.

Autonomous vehicles become idle after dropping off a passenger, and before picking up the next scheduled passenger. During this period, a vehicle needs to make a trip from the destination location of the first passenger to the origin location of the second, in a travel time window that is bounded from below by the scheduled arrival time of the first passenger, and from above by the scheduled departure time of the second. The first optimization problem ensures that this travel time window is larger than the actual travel time between the two locations. Although this time window is not strictly an idle period (the vehicle should travel to the pick-up location for the second trip), the vehicle can use the extra time to serve on-demand trip requests, and that is why we refer to the time window between two consecutive scheduled drop-off and pick-ups as the free travel time window.

In order to mathematically formulate the carsharing problem, we identify the set of free time windows between scheduled trips, and find the maximum number of carsharing requests that can be satisfied during these time windows. We keep the set of free time windows for each vehicle $v \in V$ in set $J(v)$. The e^{th} free time window of vehicle v starts after dropping off its e^{th} assigned passenger, and ends when passenger $e+1$'s pick-up has been scheduled. We denote this travel time window by $[ED_{(v,e)}\ LA_{(v,e)}]$. During this time window, the vehicle needs to travel from the destination location of its e^{th} scheduled trip to the origin location of its next scheduled trip. We denote these parameters by $OS_{(v,e)}$ and $DS_{(v,e)}$, respectively, and formulate this problem using three sets of decision variables in equations (4)-(6).

$$x_\ell^{ve} = \begin{cases} 1 & \text{If vehicle } v \text{ travels on link } \ell \text{ during its } e^{\text{th}} \text{ free time window} \\ 0 & \text{Otherwise} \end{cases} \tag{4}$$

$$y_\ell^{kve} = \begin{cases} 1 & \text{If request } k \text{ is served on link } \ell \text{ using the } e^{\text{th}} \text{ free time window of vehicle } v \\ 0 & \text{Otherwise} \end{cases} \tag{5}$$

$$z_k = \begin{cases} 1 & \text{If carsharing request } k \text{ is served} \\ 0 & \text{Otherwise} \end{cases} \tag{6}$$

Contrary to the problem in (3) where all vehicles have access to the same link set L, here each vehicle has a different link set in each of its free time windows. Let us denote by L_{ve} the set of links for vehicle v during its e^{th} free window. Furthermore, let us introduce set $L_{kve} = L_k \bigcap L_{ve}$ to contain all links that are reachable by both vehicle v during its e^{th} time window, and by request k.

The constraint sets that defines this problem are very similar to the constraint sets in the previous section, where we routed the autonomous vehicles to satisfy the set of essential trips. Constraint sets (7b)-(7d) route vehicles within their free time windows. Constraint set (7b) ensures that each vehicle at each of its free time windows leaves its origin station after delivering its last passenger. Constraint set (7c) ensures that the vehicle reaches its destination location before the departure time of its next scheduled trip. Constraint set (7d) is the flow conservation constraint. Constraint sets (7e)-(7g) route on-demand requests in the network. These sets of constraints are similar to constraint sets (7b)-(7d) that route vehicles, with a slight variation that not all on-demand requests can be necessarily served. This is reflected in the formulation by replacing the unit value on the right hand side of constraint sets (3e) and (3f) by variable z_k in constraint sets (7e) and (7g). Finally, constraint set (7h) ensures that each request is assigned to a single vehicle, and each vehicle is assigned to a single request at a time.

The objective function of the carsharing problem presented in Eq. (7a) maximizes a weighted sum of served trips. This weight can be determined based on the goal of the system. For instance, if the objective is to maximize the revenue obtained from the system, η_k can be set as the rental duration for request k. We use $\eta_k = 1$ in the experiments in this paper to maximize the total number of served requests, with the goal of finding the minimum number of vehicles required to serve the community.

$$
\begin{aligned}
\text{Maximize} \quad & \sum_{k \in M} \eta_k z_k && && \text{(7a)}\\
\text{Subject to:} \quad & \sum_{\substack{\ell \in L_{ve}:\\ s_i = OS_{(v,e)}}} x_\ell^{ve} - \sum_{\substack{\ell \in L_{ve}:\\ s_j = OS_{(v,e)}}} x_\ell^{ve} = 1 && \forall v \in V, e \in J(v) && \text{(7b)}\\
& \sum_{\substack{\ell \in L_{ve}:\\ s_j = DS_{(v,e)}}} x_\ell^{ve} - \sum_{\substack{\ell \in L_{ve}:\\ s_i = DS_{(v,e)}}} x_\ell^{ve} = 1 && \forall v \in V, e \in J(v) && \text{(7c)}\\
& \sum_{\substack{s_i \in S, t_i \in T:\\ \ell = (t_i, s_i, t, s) \in L_{ve}}} x_\ell^{ve} = \sum_{\substack{s_j \in S, t_j \in T:\\ \ell = (t, s, t_j, s_j) \in L_{ve}}} x_\ell^{ve} && \begin{array}{l} \forall v \in V, e \in J(v)\\ \forall t \in T, s \in S_e \backslash D_o :\\ \ell = (t_i, s_i, t, s) \in L_{ve} \end{array} && \text{(7d)}\\
& \sum_{\substack{v \in V\\ e \in J(v)}} \sum_{\substack{\ell \in L_{kve}:\\ s_i = OS_k}} y_\ell^{kve} - \sum_{\substack{v \in V\\ e \in J(v)}} \sum_{\substack{\ell \in L_{kve}:\\ s_j = OS_k}} y_\ell^{kve} = z_k && \forall k \in M && \text{(7e)}\\
& \sum_{\substack{v \in V\\ e \in J(v)}} \sum_{\substack{\ell \in L_{kve}:\\ s_j = DS_k}} y_\ell^{kve} - \sum_{\substack{v \in V\\ e \in J(v)}} \sum_{\substack{\ell \in L_{kve}:\\ s_i = DS_k}} y_\ell^{kve} = z_k && \forall k \in M && \text{(7f)}\\
& \sum_{\substack{s_i \in S_e, t_i \in T:\\ \ell = (t_i, s_i, t, s) \in L_{kve}}} y_\ell^{kve} = \sum_{\substack{s_j \in S_e, t_j \in T:\\ \ell = (t, s, t_j, s_j) \in L_{kve}}} y_\ell^{kve} && \begin{array}{l} \forall v \in V, e \in J(v), k \in M,\\ t \in T, s \in S_e \backslash \{OS_k, DS_k\} :\\ \ell = (t_i, s_i, t, s) \in L_{kve} \end{array} && \text{(7g)}\\
& \sum_{k \in M, \ell \in L_{kve}} y_\ell^{kve} \leq x_\ell^{ve} && \forall v \in V, e \in J(v), \ell \in L_{ve} && \text{(7h)}\\
& x_\ell^{ve}, y_\ell^{kve}, z_k \in \{0, 1\} && && \text{(7i)}
\end{aligned}
$$

4 Solution Method

We formulated the first optimization problem to find the minimum number of autonomous vehicles required to serve a cluster's set of essential trips, and optimally route these vehicles. This problem does not need be solved in real-time, and therefore for problems of moderate size (as we will discuss in the following sections) optimization engines such as CPLEX can be used to solve it.

The second optimization problem that maximizes the number of served carsharing requests may need to be solved in real-time, as carsharing requests arrive dynamically. In this section, we devise a heuristic to solve this problem in real-time. The numerical study that follows illustrates the level of efficiency and accuracy of this heuristic algorithm.

The carsharing problem as described in the previous section bears similarities to the family of parallel machine scheduling problems in manufacturing. This class of problems includes a large variety of problems, and is used to find the optimal sequence of using machinery in manufacturing processes. Parallel machine scheduling problems vary in job characteristics (whether there are preemptive or precedence constraints present, fixed/relaxed start or finish times, etc.), machine characteristics (identical or non-identical, serial or parallel, etc.), and the optimality criterion (max number of completed jobs, min makespan, etc.). In the context of our carsharing problem, jobs are carsharing requests, and machinery are the free time windows of drivers. The problem we are trying to solve has the following characteristics:

1. No preemptive or precedence constraints present: Once we fix the schedules of the essential trips, the vehicles' free time windows can be used in any manner, i.e., there is no precedence requirement on the sequence of the carsharing request to be satisfied.

2. Multiple non-homogeneous machines/servers: In our problem each free time window of each vehicle poses as a separate server. Furthermore, our servers are non-homogeneous, meaning that each vehicle at each of its free time windows has a distinct origin and destination, as well as start and finish times.

3. Jobs are available during specified time windows, rather than with specific start and finish times: the carsharing requests specify a time window during which a vehicle is required, rather than specify the exact time for start and end of their requests. Note that this does not preclude the case of requests that need the vehicle to be delivered at a specific time; for such requests, the length of the rental time window would be equal to the length of the trip for which the vehicle is being rented out.

4. Set-up cost: In our problem there exist server- and job sequence-dependent set-up costs. Because vehicles have to travel to locations where they are requested, selecting a vehicle to be assigned to a request and the sequence of requests assigned to a vehicle both play a role in the total system cost.

5. Objective: maximizing the number of served jobs (satisfied carsharing requests).

There is an extensive amount of literature on machine scheduling (Hall and Sriskandarajah 1996; Cheng et al. 2004). Rabadi et al. (2006) propose heuristics to solve the non-preemptive unrelated parallel machine scheduling problem, in which machine- and job sequence-dependent setup times are considered, but jobs are all assumed to be available at time zero. Gabrel (1995) proposes heuristics to solve the problem of scheduling non-preemptive jobs with an interval for starting time, on identical parallel machines. To the best of our knowledge, there is no study that combines both characteristics (set-up costs, and time windows for jobs), that can be used to solve the carsharing problem formulated in the previous section.

4.1 Heuristic Algorithm to Solve the on-demand Carsharing Problem

The heuristic algorithm described in this section is based on the earliest finishing time (EFT) algorithm originally designed to solve the interval scheduling problem. In the interval scheduling problem, there is a machine that needs to complete the maximum number of jobs possible. Each job has a specific start and finish time. At each step, the EFT heuristic selects the job with the earliest finishing time that does not conflict with the previously selected jobs. The EFT algorithm yields optimal solutions (Kleinberg and Tardos, 2006).

The carsharing problem is substantially more complicated than the interval scheduling problem. In fact, it is easy to see that the carsharing problem is NP-Hard. Here, we modify the EFT heuristic and tailor it to solve the carsharing problem. Our proposed algorithm is displayed in Algorithm 1.

In the mathematical program in (7), we used the tuple (v, e) to refer to the e^{th} free time window of vehicle v. In the interest of simplifying the notation, we treat each free time window of each vehicle as a separate vehicle $v' \in V'$, where $V' = \{(v, e) | v \in V, e \in J(v)\}$.

In the first step of the algorithm, we initialize two sets of arrays. The first array, $Loc(v')$, indicates the current location of vehicle v'. The second array, $Time(v')$, indicates the time vehicle v' becomes idle (available). We initialize the location array Loc for each vehicle $v' \in V'$ with the origin station of the vehicle, and the time array $Time$ with the earliest departure time of the vehicle.

Algorithm 1 On-demand vehicle allocation

Allocates carsharing requests using the idle autonomous vehicles

1. Initialize: $\forall v' \in V'$
   $Loc(v') = OS_{v'}$
   $Time(v') = ED_{v'}$
2. Find the set of feasible requests $R(v')$, $\forall v' \in V'$
   $\forall k \in R$:
   If $\quad Max\{Time(v') + shp^1(Loc(v'), OS_k), ED_k\} + shp(OS_k, DS_k) \leq LA_k$
   $R(v') = R(v') \cup k$
3. Find the matched request and driver by studying the minimum finishing time for all combinations of vehicles and requests $\big(\forall v' \in V', k \in R(v')\big)$
   $(v'^*, k^*) = Argmin_{v' \in V', k \in R(v')}\big\{Max\{Time(v') + shp(Loc(v'), OS_k), ED_k\} + shp(OS_k, DS_k)\big\}$
4. Update sets
   $Loc(v'^*) = DS_{k^*}$
   $Time(v'^*) = Max\{Time(v'^*) + shp(Loc(v'^*), OS_{k^*}), ED_{k^*}\} + shp(OS_{k^*}, DS_{k^*})$
5. Delete k^* from $R(v'), \forall v' \in V'$.
6. Update $R(v'^*)$ based on rider travel time windows:
   $\forall k \in R(v'^*)$:
   If $\quad Max\{Time(v'^*) + shp(Loc(v'^*), OS_k), ED_k\} + shp(OS_k, DS_k) > LA_k\big\}$
   $R(v'^*) = R(v'^*) \backslash k$
   Go to step 3.
7. Stopping Criteria: $\forall v' \in V', R(v') = \emptyset$.

[1] $shp(i, j)$: shortest path travel time between i and j

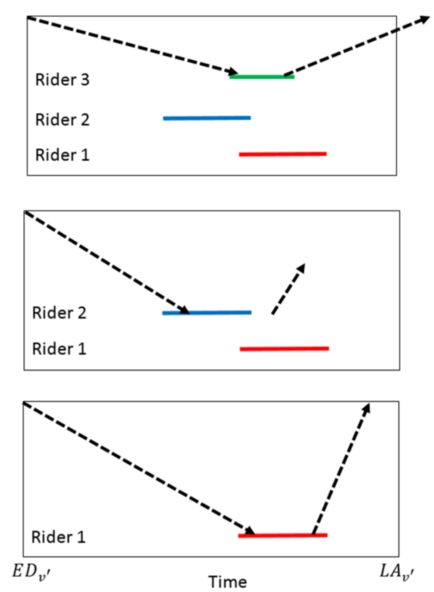

Figure 3: Determining the set of feasible requests $R(v')$ for vehicle v'

The algorithm starts by determining the set of feasible carsharing requests for each vehicle. In order for a request to be feasible for a vehicle, the vehicle should be able to drive from its current location to the request's origin, stay in possession of the renter for the requested duration of time, and finally arrive at its own destination (the pick-up location of its next scheduled essential trip) before its latest arrival time.

Figure 3(b) studies the feasibility of three carsharing requests for an example vehicle. The boundaries of the boxes show the free time window of the vehicle, and the line (marked by colors blue, red, or green) associated with each request marks its time window. The first request (at the bottom) is feasible for the vehicle: the vehicle arrives at the request's origin location after the request's earliest departure time, is able to stay in possession of the renter for the requested duration (that ends before the request's latest arrival time), and travels to its destination within its time window. The second and third requests, however, are not feasible for the vehicle. In case of the second request, the vehicle cannot stay in possession of the renter for the duration of the request, and in case of the third request, the vehicle cannot arrive at its own destination in time, after finishing serving the request.

In the third step the algorithm finds finishing times for all combinations of vehicles and their set of feasible requests. The finishing time for vehicle v serving request k is sum of the time required for the vehicle to arrive at the request's origin, the rental period, and the waiting time for the rental window to start in case the vehicle arrives at the rental location in advance of the rental window.

The vehicle and request pair that lead to the earliest finishing time will be selected and matched together. In step 4, the location of the matched vehicle will be updated to the destination of the matched request, and the time array of the assigned vehicle will be updated to the drop-off time of the rented out vehicle. In step 5, the matched request in step 3 will be eliminated from the set of available requests to all vehicles. Furthermore, since the time window and location of the matched vehicle have been updated, the set of feasible requests for this vehicle needs to be updated as well. The algorithm stops when sets of feasible requests for all vehicles become empty.

5 Experiments

We implemented the SVOR program for a sample of households in San Diego County, using data from the 2000-2001 California statewide household travel survey (Casas, 2002). In this survey, Caltrans collected travel data from 17049 volunteer households in California. These households were selected carefully to ensure that the sample would be a good representation of the state population. After cleaning the data by eliminating records with incomplete or contradicting information, we identified a total of 1184 households residing in the city of San Diego. For these households, detailed information on the number of household members, number of vehicles owned by households, and logged trips for each household member during a working day along with the purpose of each trip were available, among other information.

We determined the set of essential trips for each household based on the information on the purpose of trips. All car-based trips (e.g., private cars, public and private shuttles, public transportation) reported by households were considered in this study, regardless of the reported mode. Note that over 70% of the initial modal share collected through the survey belongs to private vehicles, while 25% of the modal share is unspecified, 3% belongs to walking, and the remaining 2% is collectively claimed by bus, heavy and light rail, and bike. We categorized trips concerning work, school, childcare, medical, fitness, community meetings, volunteer activities, visiting friends and family, and entertainment activities as essential (core), and the rest of the trips as non-essential. After cleaning the dataset, among the 1184 households, 573 of them were left without any essential trips during the survey day, and therefore were not considered for the shared use and ownership program. These households, however, were taken into consideration for the car rental service to serve their non-essential trips. The 1184 households made a total of 3306 trips, 1624 (49%) of which were essential trips.

The first step in implementing the program is to cluster households. Each cluster should include a number of households with enough commonalities that would interest them to participate in the shared vehicle ownership and ridership program together. Various parameters can be used to determine a suitable cluster for a household, including home location, demographics and socio-economic status of households, level of spatiotemporal proximity of trips between households, and income level, to name a few. In this study, we cluster households at different levels and using different criteria. In sections 5.1 and 5.2 we study the impact of clustering households based on the proximity of their residence locations and the degree of overlap between their trips, respectively. In section 5.3, we study two extreme clustering approaches and estimate upper and lower bounds on the potential savings on the number of vehicles and VMT in our sample.

5.1 Residence-based Clustering

In our first implementation, we use agglomorative clustering, an unsupervised learning method, to group households based on their home location (Steinbach et al., 2000). Figure 4 displays the resulting 277 clusters of households. These clusters are distinguished by color based on their size (i.e., the number of their household members). Figure 4 also displays the distribution of number of households in clusters. About 80% of clusters have three or fewer household members, which makes it an easier task to manage shared vehicles when it comes to incorporating the preferences of members on whom to add to the clusters in the future, or the cluster members with whom to share rides. About 30% of the households are geographically isolated from others and therefore remain as singleton clusters. Figure 4 also illustrates the Voronoi polygons attributed to clusters. These polygons suggest to which cluster a prospective household looking to join the program would belong based on its residence location.

For each cluster, we solve the optimization problem (3) to find the optimal number of vehicles required to serve the cluster's entire set of essential trips. All problems are solved on a PC with Core i7 3GHz and 8GB of RAM, using the AMPL modeling language and CPLEX 12.6.00 solver with standard tuning. Solution times are displayed in Figure 5. Not surprisingly, solution times increase with cluster size; however, they remain within a reasonable range for a problem that does not need to be solved in real-time. Figure 5 suggests that the mathematical formulation in problem (3) is scalable with respect to the number of households in clusters.

The solution suggests that a total of 379 vehicles are required to serve all the essential trips by all clusters (including the single household clusters). Note that households are still in need of transportation for their non-essential trips. Therefore, this number serves as a lower bound to the total number of required vehicles.

Figure 6(a) shows the distribution of the number of vehicles owned by households in year 2000. This figure suggests that the majority of households owned at least 2 vehicles. Figure 6(b) shows the distribution of the optimal number of vehicles needed under the SVOR program. This figure shows that about 65% of clusters need no more than one vehicle to serve their essential trips. No cluster needs more than 4 vehicles.

After forming clusters of households and determining the fixed schedules of autonomous vehicles to serve their cluster members, we need to address the non-essential trips. One possibility is to use a central carsharing system that rents out the autonomous vehicles owned by clusters to serve non-essential trips of the entire population. The question is, what percentage of the non-essential trips

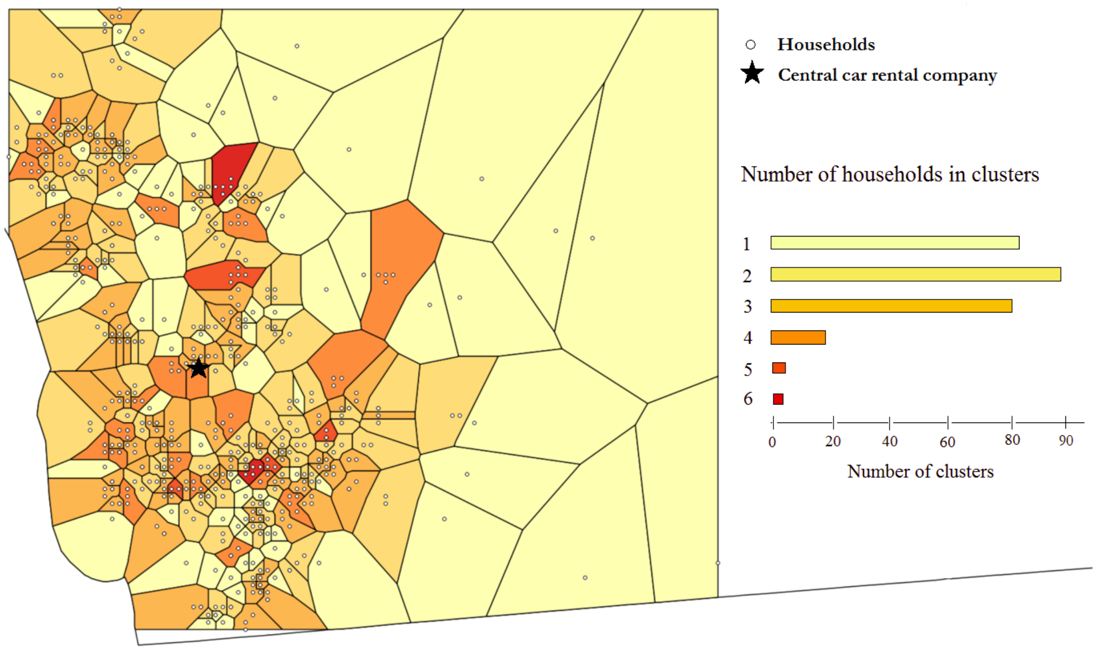

Figure 4: Clusters of households. Households in each cluster are assumed to share ownership of a set of autonomous vehicles

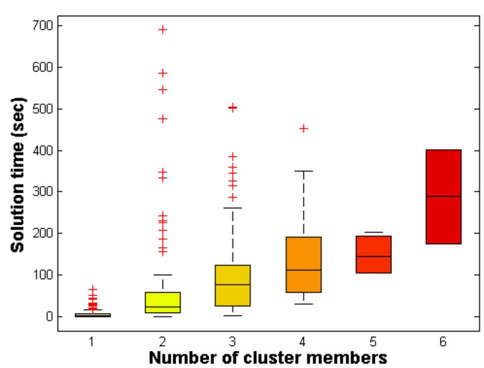

Figure 5: Solution time (sec) of finding the optimal number of vehicles and vehicle itineraries for each cluster size

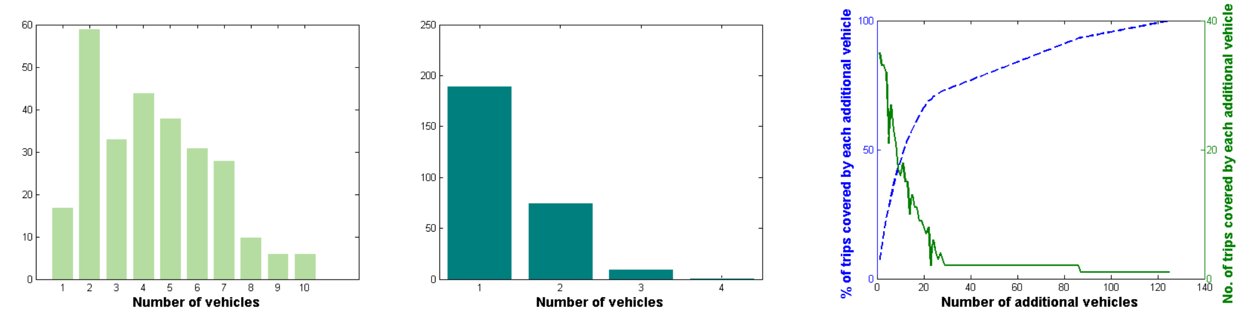

(a) Frequency distribution of the number of vehicles owned by cluster members in year 2000

(b) Frequency distribution of the number of autonomous vehicles required for clusters under the SVOR program

(c) Number of additional vehicles required to serve non-essential trips

Figure 6: Impact of the SVOR program on vehicle ownership under residence-based clustering

can be served by such a centralized system, and how many additional vehicles need to be owned by the carsharing system to serve the entire set of non-essential trips of the population. Note that not only the 611 households who participate in the SVOR program, but also the 573 households who did not report any essential trips need to have their non-essential trips served.

Using Algorithm 1 to rent out the 379 autonomous vehicles owned by clusters with the goal of serving as many non-essential trips as possible, we manage to serve 63% of the non-essential trips. We then use a variation of Algorithm 1 to find the additional number of vehicles the rental company needs to own in order to serve the remaining 37% of the non-essential trips. We assume a depot for the rental service provider located strategically in the network (marked with a star symbol in Figure 4) where there is a very high density of trip origin and destination locations. We increment the number of vehicles one at a time, having each vehicle start its trip from the depot, to which it returns at the end of the day. We assign trips to each newly added vehicle using Algorithm 1 until all trips are served.

The results suggest that a total of 125 vehicles need to be owned and managed by the rental service provider to serve the rest of the non-essential trips. All these vehicles, however, do not have the same contribution in terms of the number of served trips. Figure 6(c) displays the cumulative percentage and number of trips served by each additional vehicle. This figure suggests that using only 30 vehicles, the rental service provider can serve 75% of the remaining non-essential trips. The remaining 95 vehicles each serve only 1 or 2 trips with origins and/or destinations in remote areas. In fact, using cluster-owned vehicles in their idle times, and owning an additional 30 vehicles, the rental company can serve more than 90% of the on-demand trips. Comparing the count of 504 vehicles that could serve the entire transportation demand against the 2194 vehicles owned by households suggests that the SVOR program has the potential to have a significant impact on vehicle ownership.

The savings in the number of vehicles in the proposed system originate from three different sources: (1) introduction of autonomous vehicles, (2) shared ownership of these vehicles, and (3) ridesharing within clusters. It would be interesting to observe how much of the savings can be attributed to each source. Towards this end, we consider two additional cases. In the first case, we study the impact of households trading their current vehicles for the optimal number of autonomous vehicles. In the second case, we allow households to form clusters in order to share the ownership of vehicles. The third and most comprehensive case considers the shared ownership and shared

Table 1: Impact of different elements of the SVOR program on vehicle ownership and vehicles miles traveled using household locations as clustering criterion

	Base Case: Year 2000 data	Case 1: Autonomous vehicles	Case 2: Autonomous vehicles + Shared ownership	Case 3: Autonomous vehicles + Shared ownership + Shared ridership
No. of vehicles	2,194	787	528	504
Total VMT	29,959	62,980	64,836	65,610
Average VMT	13.6	80	122	130

ridership of vehicles among the members of each cluster, as described by the SVOR program.

Table 1 summarizes the number of required vehicles, and the total and average VMT for each case. Number of vehicles determines the initial investment, and average VMT per vehicle translates into life expectancy and operational costs. While these two costs may directly impact the cost of vehicle ownership, total VMT translates into the environmental cost of switching to autonomous vehicles. Note that for now we assume that there is parking space available to vehicles when and where required.

In case 1, where legacy vehicles are replaced with autonomous vehicles, a total of 787 vehicles need to be purchased by the households in our sample. Not surprisingly, replacing legacy vehicles with autonomous vehicles results in an increase in total and average VMT, since the number of autonomous vehicles is substantially less than the number of legacy vehicles, implying that autonomous vehicles have to make empty trips in order to serve their owners. In the second case, allowing shared ownership of autonomous vehicles leads to a 30% reduction in vehicle ownership compared to case 1, although this reduction in the number of vehicles comes at the price of a 50% increase in average VMT for each vehicle. The increase in total VMT, however, is not substantial. Finally, in the third case where SVOR of autonomous vehicles is studied, an additional 5% decline in vehicle ownership compared to case 2 can be witnessed, accompanied by a slight increase in both total and average VMT.

Note that the number of vehicles needed in the SVOR program only provides a lower bound, since trips that do not happen regularly might not have been occurred during the survey day. Additionally, the convenience of travel offered by autonomous vehicles can induce new trips. Part of these trips that have not been accounted for can be served using the unutilized capacity in the current system. For others, however, additional vehicles might be required. The number of vehicles under the 3 cases in Table 1, however, provide a solid basis for quantifying the effect of each element of the SVOR program.

5.2 Clustering Based on Trip Overlap

An alternative way of clustering households is based on the overlap between household trips. For each household pair, we compute the *degree of compatibility* between their trips. For a given pair of households h_1 and h_2, we define an $n_1 \times n_2$ matrix, where n_1 and n_2 denote the number of essential trips by the two households, respectively. Each cell c_{ij} in this matrix assumes the value 1 if trips i and j can be fulfilled using the same vehicle, i.e., one of the two following conditions holds: (i) the vehicle can fulfill the trips sequentially (i.e., fulfill one trip, followed by the second), and (ii), the vehicle can fulfill the trips concurrently (i.e., perform the pick-up task for both trips (in any sequence), followed by the drop-off task for the two trips (in any sequence)). If neither of these two conditions is satisfied, the cell c_{ij} will assume the value 0, implying that the two trips are incompatible. The summation of all elements in matrix C is what we define as the degree of

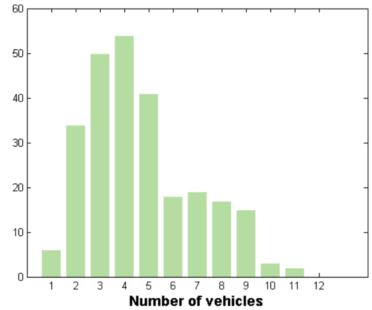

(a) Histogram of number of vehicles currently owned by cluster members

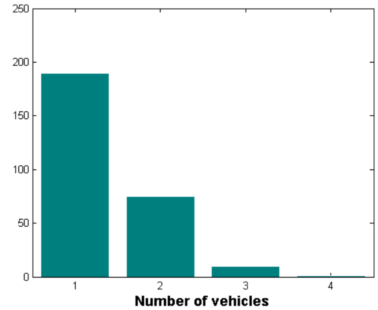

(b) Histogram of optimized number of autonomous vehicles required for clusters

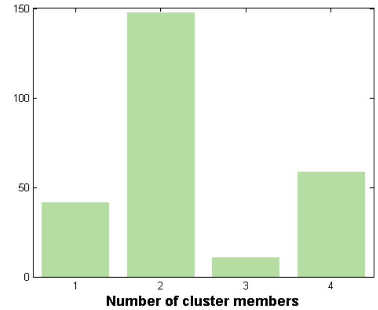

(c) Histogram of optimized number of autonomous vehicles required for clusters

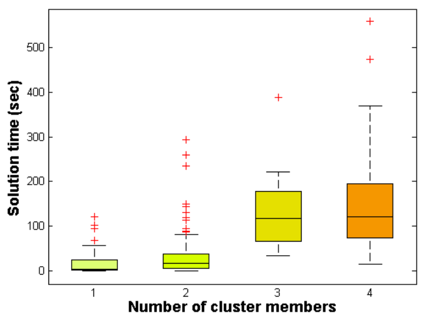

(d) Solution time (sec) of finding the optimal number of vehicles and vehicle itineraries for each cluster size

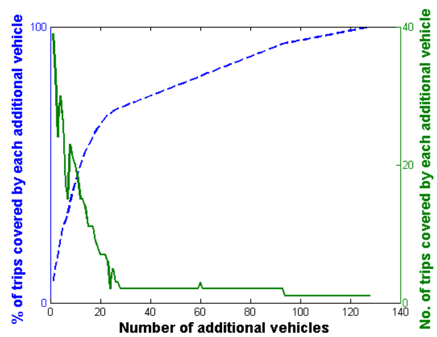

(e) Number of additional vehicles required to serve the entire set of non-essential trips

Figure 7: Impact of the SVOR program on vehicle ownership under clustering based on degree of inter-household trip overlaps

compatibility between households h_1 and h_2. We use agglomerative clustering based on the degree of compatibility between households to group households into 261 clusters.

A total of 340 autonomous vehicles are required to serve the essential trips of the 261 clusters. Figures 7(a) and 7(b) compare the distribution of vehicles owned by households before and after implementing the SVOR program. It is interesting to note that more than half of the clusters need only one autonomous vehicle, whereas in the base case the majority of clusters owned four or more vehicles (Numbers for the base case are obtained by adding the number of vehicles owned by all households in a cluster based on the survey data). Figures 7(c) and 7(d) display the distribution of cluster size, and the solution times to optimize the itineraries in clusters, respectively. Similar to the residence-based clustering, the solution times remain within a reasonable range.

During their idle times, essential vehicles managed by the rental service provider can serve 57% of the non-essential trips of the entire population. The company needs to own 128 additional vehicles to serve the remaining 43% of the non-essential trips. Figure 7(e) shows the contribution of each additional vehicle to serving the remainder of non-essential trips. This figure suggests that although 128 vehicles are required to serve the entire set of non-essential trips, the first 27 vehicles can serve more than 70% of the remaining set of non-essential trips, and along with the cluster vehicles, around 90% of the entire set of non-essential trips.

Similar to the previous section, we study the impact of each component of the SVOR program

Table 2: Impact of different elements of the SVOR program on vehicle ownership and vehicles miles traveled using degree of trip overlaps as clustering criterion

	Base Case: Year 2000 data	Case 1: Autonomous vehicles	Case 2: Autonomous vehicles + Shared ownership	Case 3: Autonomous vehicles + Shared ownership + Shared ridership
No. of vehicles	2,194	787	497	468
Total VMT	29,959	62,980	102,790	73,841
Average VMT	13.6	80	207	158

on vehicle ownership and VMT. Results are displayed in table 2. The same declining trend in vehicle ownership that was observed in residence-based clustering can be witnessed in trip-overlap clustering as well, as we move from case 1 to case 3. For each case, however, the total number of vehicles in trip-overlap clustering is less than that of residence-based clustering. This is an intuitive finding, since in trip-overlap clustering we are grouping households whose trips have higher degrees of compatibility and therefore can be served using fewer vehicles. Fewer vehicles in this case translates into higher total and average VMT. One interesting observation is that in case 3 where we allow households in a cluster to rideshare, the total and average VMT are less than case 2 in which ridesharing is not allowed. While in residence-based clustering (Table 1) adding the ridesharing capability slightly increased the total and average VMT due to lower number of vehicles, this is not the case in trip-overlap clustering, since the clustering criteria is partially based on the compatibility of households to rideshare.

5.3 Level of Clustering

In sections 5.1 and 5.2 we studied the impact of clustering households based on two different criteria. In this section, we present two additional and more extreme clustering approaches to provide some bounds on the impact of shared vehicle ownership and ridership program for our sample of households. Results are displayed in table 3.

The first extreme approach is to conduct no clustering at all, and assume each household as a singleton cluster, resulting in 611 clusters. Our analysis shows that these households need 678 vehicles to cover their essential trips. To satisfy the non-essential trips of all 1184 households in our data set, The car rental service provider needs to own 109 separate vehicles in addition to having access to household-owned vehicles during their idle times.

By simply replacing legacy vehicles with the optimal number of autonomous vehicles, the total number of vehicles owned by the entire sample reduces from 2194 to 787. This 2.8 fold reduction in vehicle ownership, however, comes with the side effect of higher VMT for vehicles, as discussed in the previous section. The amount of increase in VMT, however, depends highly on the availability of parking. Here, we discuss two extreme scenarios on parking availability. In the first scenario, parking is available when and where required, and in the second scenario, there is no parking available, i.e., vehicles may have to drive around. In the latter case, if time allows, vehicles can travel to the closest available parking spot owned by any of its cluster households. Unavailability of parking is not limited to physical availability, but the affordability of the available parking spaces as well. If parking is expensive, it might be a financially wiser alternative for vehicles to drive around.

Our analysis shows that availability of affordable parking opportunities has a substantial impact on the total VMT. For the households in our sample, we find a 5-fold increase in the total and average VMT when there is not access to affordable parking. This result, however, is obtained under the assumption that introduction of autonomous vehicles does not change household travel

Table 3: Impact of level of clustering and parking availability of vehicle ownership and VMT

Clustering method	Household-based (no clustering)	Residence -based	Trip overlap	Single cluster
No. of clusters	611	277	261	1
Avg. cluster size	1	2.21	2.34	1184
Min no. of autonomous vehicles owned by clusters	678	379	340	258
Percentage of non-essential trips covered	63%	63%	57%	100%
Additional no. of vehicles required	109	125	128	0
Total number of vehicles	787	504	468	258
Parking not available:				
Optimized total VMT	335,640	305,800	286,060	158,020
Current avg. VMT	13.6	13.6	13.6	13.6
Optimized avg. VMT/shared vehicle	450	715	734	612
Optimized avg. VMT/additional vehicle	278	278	286	0
Optimized avg. VMT/vehicle	426	607	611	612
Parking available when and where desired:				
Optimized total VMT	62,980	65,610	73,841	76,134
Optimized avg. VMT/shared vehicle	73	134	174	295
Optimized avg. VMT/additional vehicle	121	119	115	0
Optimized avg. VMT/vehicle	80	130	158	295

patterns.

In the second extreme case, we group all households into a single cluster. In this case, a total of 258 vehicles would suffice to serve all trip requests, which is a 3-fold reduction compared to the first extreme case, where each household formed a singleton cluster. An interesting observation is that even under the assumption of unavailability of affordable parking, this case leads to substantial savings in the total VMT. This observation has two implications. First, as the number of households participating in the program and hence the average cluster size increases, the impact of parking availability becomes less prominent. The reason for this trend is that at higher levels of demand, vehicle trips would be booked back to back, not leaving time for vehicles to pass by aimlessly driving. Second, the substantial reduction in VMT along with a decline in the number of vehicles observed under the single-cluster scenario demonstrate the significant benefits that the merger of ridesharing services and autonomous vehicles can provide in transportation systems.

The two clustering approaches from sections 5.1 and 5.2 lie between the two extreme cases studied above. Although the difference between the residence-based and trip overlap-based clustering is not substantial, the results indicate that the benefits of clustering households based on trip overlap would become more prominent with higher participation rates.

Among the clustering approaches discussed above, the last scenario which considers a single cluster renders the most benefits. In this scenario, a regional shared-mobility service provider owns and manages a fleet of autonomous vehicles that serves the entire transportation needs of a region. We can, in fact, expect this type of shared-mobility service to emerge before autonomous vehicles are offered for private ownership, for two reasons. First, in order for autonomous vehicles to be ready for consumer use, high-definition maps of the entire United States road network need to be developed. The type of maps required for this purpose need to include far more details than currently contained in street maps maintained by companies like Alphabet (i.e., the Google Maps). To have fleets of autonomous vehicles operate in certain urban regions would require far less effort in collecting the network data on the targeted areas. Therefore, for areas with higher levels of demand creating high-definition maps can be prioritized, given the time-consuming nature of creating these maps. Additionally, investments in developing high-definition maps would be more economical for highly populated areas, and crowd-sourced information will be available more abundantly to keep these maps updated. A centralized service provider is more likely to be able to make the investment required for creating high definition maps, at least for regions with high levels of demand, while this is less likely to be the case for each individual household.

Second, a fleet of autonomous vehicles used in a shared-mobility context in populous areas is more likely to have a higher utilization rate, which could amount to a possibly lower cost to the consumer. High demand makes such shared-mobility services less sensitive to the high cost of autonomous technology, and the inexpensive transportation alternative provided to the consumers may render vehicle ownership in certain urban regions obsolete.

This type of regional shared-mobility service, however, might not be accessible or attractive to everyone. It might not be a financially feasible strategy for shared-mobility service providers to target households who reside in remote areas. In addition, there will always be individuals who would like to ensure their privacy, flexibility, or timeliness by owning private vehicles. Such individuals could seek other clustering approaches, or at extreme cases, form individual clusters.

6 Discussion

Throughout this paper, we made the assumption that the introduction of autonomous vehicles does not impact household travel behavior. However, autonomous vehicles can change household travel patterns in multiple ways. Having access to a new technology that allows individuals to make use of their time while traveling can encourage longer trips that would have otherwise been avoided due to the burden of driving. Moreover, access to autonomous vehicles could induce higher number of trips. With self-driving vehicles, trip chaining (which is currently a necessity to many households) would not be as essential. Self-driving cars can transport household members without a valid driver's license and perform activities such as parking or refueling in their idle times, making current trip chains smaller, changing travel patterns, and increasing the number of trips. Longer and more frequent trips impose higher costs on the transportation infrastructure, and can cancel out some of the benefits that autonomous cars introduce by reducing the number of vehicles.

The results of our study suggests that availability of affordable parking is a major determinant of the total VMT and the consequent cost to the transportation infrastructure and the environment. We demonstrated, however, that this impact depends on the degree of resource-sharing. With the right type of clustering and larger cluster sizes, total VMT may stay within a reasonable range. In the best case scenario where a central shared-mobility service provider serves the entire transportation demand of a region, there is only a 20% increase in total VMT, compared to the household-based clustering scenario where households simply replace their legacy vehicles with autonomous ones.

It should be noted the analysis presented in this paper do not capture all the effects of introducing autonomous vehicles. In spite of a possible increase in VMT, autonomous vehicles can significantly reduce the adverse impact of transportation on the environment by targeting two main sources of emission dissemination. Autonomous vehicles can reduce the number of vehicles required by orders of magnitude, as demonstrated in this study, reducing part of the congestion created as a result of higher VMT, under the right policies (e.g., VMT-based tax). In addition, these vehicles significantly increase network capacity, as we will discuss next. The combination of these two factors can reduce, and in some regions completely eliminate, the stop-and-go conditions that are a major contributor to vehicle emissions.

In addition to fewer vehicles, autonomous vehicles can further contribute to congestion-relief due to their ability to communicate. Tientrakool (2011) shows that a highway populated with a mix of autonomous and legacy vehicles can experience a substantial increase in capacity depending on the percentage of communicating vehicles in the traffic mix. Her study suggests that when less than 30% of the traffic mix can communicate, the resulting rate of change of capacity is rather slow. With penetration rate of $30\% - 85\%$ the rate of change of improvement in capacity increases, and when the percentage of communicating vehicles in the traffic mix exceeds 85%, the resulting rate of change of increase in capacity improves very quickly, to the extent that at the 100% penetration rate, the capacity reaches more than $10,000$ vehicles/hours/lane, to the point of reaching the capacity of more than $10,000$ vehicles/hours/lane under full penetration. The increase in capacity is caused by lower inter-vehicle gaps that need to be maintained (Schakel et al., 2010; Levin and Boyles, 2016). Furthermore, the communication capabilities can reduce the probability of traffic break-downs, hence maintaining higher capacities under non-idealized conditions (Kesting et al., 2010).

Autonomous vehicles eliminate the possibility of human error, which is the leading cause for the majority of traffic collision fatalities. An autonomous driver will never get distracted, fall asleep, or drive under the influence. Furthermore, autonomous vehicles can make split-second decisions based on probabilistic models fed by far more complete information than a human driver can have access to, while a human driver needs to take longer to make a decision based on in-complete information. As the percentage of autonomous vehicles in the traffic mix increases, so does the level of safety, as the network becomes more deterministic.

Although beyond the scope of this paper, the complicated nature of actual deployment of autonomous vehicles goes beyond assessing their level of contribution in enhancing safety, mobility, and the environment. Legal liability in car collisions (Duffy and Hopkins, 2013), insurance matters (Peterson, 2012), and efficient management of shared fleet, which all become even more complicated in the context of a shared ownership and ridership model, are all matters that should be studied.

To conclude, although change in travel patterns and increase in VMT under certain conditions may lead to higher costs to the transportation system, the many benefits of autonomous vehicles described in this paper may more than outweigh the possible downfalls, if the right policies are put in place. Furthermore, higher contribution rates in the SVOR program would exponentially increase the benefits. Note that the sample we used for this study contained only 0.1% of the population, which is far less than the participation rate expected for such services. Despite the small sample size, we showed that in the case of a central shared-mobility service provider we can achieve a 9-fold reduction in the number of vehicles. Although these savings fall within the previously suggested range of 6-10 fold reduction in the number of vehicles, note that these savings can increase with the penetration rate of shared ridership. Much of the the higher success rate of the SVOR program proposes in this paper can be attributed to its "shared ridership" component.

Finally, introduction of autonomous vehicles in the traffic mix may call for some policy adjustments to reduce the possible adverse impacts of this new technology. Our study points out that parking availability and cost play important roles on the potential environmental and congestion-relief benefits expected from autonomous vehicles. A grid-lock for an autonomous vehicle that is traveling aimlessly is equivalent to free parking. The temptation to drive around when idle may become more strong with alternative-fuel autonomous vehicles, since under current regulations, roads and highways are funded by gasoline tax, and therefore the only driving related cost alternative-fuel autonomous vehicles would have to bear is the depreciation cost associated with higher VMT. To avoid such behavior, changing from gasoline-based to VMT-based tax might be a necessary policy adjustment.

7 Conclusion

In this paper, we proposed a mathematical framework to model shared ownership and ridership of autonomous vehicles. The motivation behind this model is to assess the impact of switching to autonomous vehicles on the number of vehicles and total VMT, under various degrees of shared ownership. For a group of households willing to participate in the program together, we formulated an optimization problem to find the minimum number of vehicles required to satisfy their transportation needs. The program allows participants to register their vehicles in a central carsharing program when they were not being used, in order to generate revenue. We implemented this program for a sample of households in the San Diego County, California, and studied the impact of different clustering criteria on vehicle ownership and vehicle miles traveled.

Our study suggests that self-driving vehicles, when used in a shared setting, can introduce a much more flexible and inexpensive form of shared-mobility in certain populous regions, rendering vehicle ownership and public transit in its current form obsolete. For other areas with less dense populations, replacing legacy vehicles with autonomous cars, especially under a shared ownership program, can still introduce benefits in terms of reduction in vehicle ownership. Our study also suggests that the extent of environmental and congestion-relief benefits expected from autonomous vehicles depends on operational and deployment strategies such as availability of affordable parking.

References

Barth, M. and S. Shaheen (2002). Shared-use vehicle systems: Framework for classifying carsharing, station cars, and combined approaches. *Transportation Research Record: Journal of the Transportation Research Board* (1791), 105–112.

Casas, J. (2002). 2001 california statewide household travel survey final report. *California Department of Transportation*.

Cervero, R., A. Golub, and B. Nee (2007). City carshare: longer-term travel demand and car ownership impacts. *Transportation Research Record: Journal of the Transportation Research Board* (1992), 70–80.

Cervero, R. and Y. Tsai (2004). City carshare in san francisco, california: second-year travel demand and car ownership impacts. *Transportation Research Record: Journal of the Transportation Research Board* (1887), 117–127.

Chan, N. D. and S. A. Shaheen (2012). Ridesharing in north america: Past, present, and future. *Transport Reviews 32*(1), 93–112.

Cheng, T. E., Q. Ding, and B. M. Lin (2004). A concise survey of scheduling with time-dependent processing times. *European Journal of Operational Research 152*(1), 1–13.

Cordeau, J.-F. and G. Laporte (2007). The dial-a-ride problem: models and algorithms. *Annals of Operations Research 153*(1), 29–46.

Duffy, S. and J. P. Hopkins (2013). Sit, stay, drive: The future of autonomous car liability.

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

Fagnant, D. J. and K. M. Kockelman (2015a). Dynamic ride-sharing and optimal fleet sizing for a system of shared autonomous vehicles. In *Transportation Research Board 94th Annual Meeting*, Number 15-1962.

Fagnant, D. J., K. M. Kockelman, and P. Bansal (2015b). Operations of shared autonomous vehicle fleet for Austin, Texas, market. *Transportation Research Record: Journal of the Transportation Research Board* (2536), 98–106.

Gabrel, V. (1995). Scheduling jobs within time windows on identical parallel machines: New model and algorithms. *European Journal of Operational Research 83*(2), 320–329.

Gao, Z., J. Wu, and H. Sun (2005). Solution algorithm for the bi-level discrete network design problem. *Transportation Research Part B: Methodological 39*(6), 479–495.

Garigliano, J. (2007). Own a lamborghini-part time. Accessed July 22, 2015.

Hall, N. G. and C. Sriskandarajah (1996). A survey of machine scheduling problems with blocking and no-wait in process. *Operations research 44*(3), 510–525.

Heinrich, S. (2010). *Implementing real-time ridesharing in the San Francisco Bay Area.* Ph. D. thesis, San Jose State University.

Jaw, J.-J., A. R. Odoni, H. N. Psaraftis, and N. H. Wilson (1986). A heuristic algorithm for the multi-vehicle advance request dial-a-ride problem with time windows. *Transportation Research Part B: Methodological 20*(3), 243–257.

John, S. and S. Phadnis (2015). Ford to pilot car-sharing project in Bengaluru. Accessed March 10, 2017.

Kesting, A., M. Treiber, and D. Helbing (2010). Enhanced intelligent driver model to access the impact of driving strategies on traffic capacity. *Philosophical Transactions of the Royal Society of London A: Mathematical, Physical and Engineering Sciences 368*(1928), 4585–4605.

Kleinberg, J. and E. Tardos (2006). *Algorithm design.* Pearson Education India.

Krueger, R., T. H. Rashidi, and J. M. Rose (2016). Preferences for shared autonomous vehicles. *Transportation research part C: emerging technologies 69*, 343–355.

Levin, M. W. and S. D. Boyles (2016). A multiclass cell transmission model for shared human and autonomous vehicle roads. *Transportation Research Part C: Emerging Technologies 62*, 103–116.

Martin, E., S. Shaheen, and J. Lidicker (2010). Impact of carsharing on household vehicle holdings: Results from north american shared-use vehicle survey. *Transportation Research Record: Journal of the Transportation Research Board* (2143), 150–158.

Masoud, N. and R. Jayakrishnan (2017). A decomposition algorithm to solve the multi-hop peer-to-peer ride-matching problem. *Transportation Research Part B: Methodological 99*, 1–29.

Naughton, K. (2015). Driverless, shared cars to cut vehicle ownership by half: Barclays. http://www.insurancejournal.com/news/national/2015/05/19/368711.htm. Accessed Oct. 6, 2015.

Peterson, R. W. (2012). New technology-old law: Autonomous vehicles and California's insurance framework. *Santa Clara L. Rev. 52*, 1341.

Prasad, P. (Aug. 2016). Ford leads 24 million dollar investment in self-drive car rental startup Zoomcar. Accessed March 10, 2017.

Psaraftis, H. N. (1983). An exact algorithm for the single vehicle many-to-many dial-a-ride problem with time windows. *Transportation Science 17*(3), 351–357.

Rabadi, G., R. J. Moraga, and A. Al-Salem (2006). Heuristics for the unrelated parallel machine scheduling problem with setup times. *Journal of Intelligent Manufacturing 17*(1), 85–97.

Schakel, W. J., B. Van Arem, and B. D. Netten (2010). Effects of cooperative adaptive cruise control on traffic flow stability. In *Intelligent Transportation Systems (ITSC), 2010 13th International IEEE Conference on*, pp. 759–764.

Schall, E. (2015). The future of car ownership and autonomous driving. http://www.cheatsheet.com/automobiles/the-future-of-car-ownership-and-autonomous-driving.html/?a=viewall. Accessed Aug. 1 2015.

Schoettle, B. and M. Sivak (2015). Potential impact of self-driving vehicles on household vehicle demand and usage. *University of Michigan Transportation Research Institute*.

Steinbach, M., G. Karypis, V. Kumar, et al. (2000). A comparison of document clustering techniques. In *KDD workshop on text mining*, Volume 400, pp. 525–526.

Tientrakool, P. (2011). *Reliable Neighborcast Protocol for Vehicular Ad hoc Networks.* Ph. D. thesis, Columbia University.

Walb, C., W. Loudon, and C. Systematics (1986). *Evaluation of the Short-Term Auto Rental (STAR) Service in San Francisco, CA.* US Department of Transportation, Urban Mass Transportation Administration.

Wickerham, J. (Mar. 2017). What it really costs to turn a car into a self-driving vehicle. https://qz.com/924212/what-it-really-costs-to-turn-a-car-into-a-self-driving-vehicle/, Accessed May 10, 2017.