International Transport Forum

CPB
Corporate Partnership Board

Urban Mobility System Upgrade

How shared self-driving cars could change city traffic

Corporate Partnership Board Report

OECD

Urban Mobility System Upgrade

How shared self-driving cars could change city traffic

About the International Transport Forum

The International Transport Forum at the OECD is an intergovernmental organisation with 54 member countries. It acts as a think tank for transport policy and organises the Annual Summit of transport ministers. ITF is the only global body that covers all transport modes.

ITF works for transport policies that improve peoples' lives. Our mission is to foster a deeper understanding of the role of transport in economic growth, environmental sustainability and social inclusion and to raise the public profile of transport policy.

ITF organises global dialogue for better transport. We act as a platform for discussion and pre-negotiation of policy issues across all transport modes. We analyse trends, share knowledge and promote exchange among transport decision makers and civil society.

Our member countries are: Albania, Armenia, Australia, Austria, Azerbaijan, Belarus, Belgium, Bosnia and Herzegovina, Bulgaria, Canada, Chile, China (People's Republic of), Croatia, Czech Republic, Denmark, Estonia, Finland, France, Former Yugoslav Republic of Macedonia, Georgia, Germany, Greece, Hungary, Iceland, India, Ireland, Italy, Japan, Korea, Latvia, Liechtenstein, Lithuania, Luxembourg, Malta, Mexico, Republic of Moldova, Montenegro, Netherlands, New Zealand, Norway, Poland, Portugal, Romania, Russian Federation, Serbia, Slovak Republic, Slovenia, Spain, Sweden, Switzerland, Turkey, Ukraine, United Kingdom and United States.

www.internationaltransportforum.org

Table of contents

Tables

Figures

Executive summary

This report examines the changes that might result from the large-scale uptake of a shared and self-driving fleet of vehicles in a mid-sized European city. The study explores two different self-driving vehicle concepts, for which we have coined the terms "TaxiBot" and "AutoVot". TaxiBots are self-driving cars that can be shared simultaneously by several passengers. AutoVots pick-up and drop-off single passengers sequentially. We had two premises for this study: First, the urban mobility system upgrade with a fleet of TaxiBots and AutoVots should deliver the same trips as today in terms of origin, destination and timing. Second, it should also replace all car and bus trips. The report looks at impacts on car fleet size, volume of travel and parking requirements over two different time scales: a 24-hour average and for peak hours only.

What we found

Nearly the same mobility can be delivered with 10% of the cars

TaxiBots combined with high-capacity public transport could remove 9 out of every 10 cars in a mid-sized European city. Even in the scenario that least reduces the number of cars (AutoVots without high-capacity public transport), nearly eight out of ten cars could be removed.

The overall volume of car travel will likely increase

A TaxiBot system with high-capacity public transport will result in 6% more car-kilometres travelled than today, because these services would have to replace not only those provided by private cars and traditional taxis but also all those provided by buses. An AutoVot system in the absence of high-capacity public transport will nearly double (+89%) car-kilometres travelled. This is due to repositioning and servicing trips that would otherwise have been carried out by public transport.

Impacts on congestion depend on system configuration

A TaxiBot system in combination with high-capacity public transport uses 65% fewer vehicles during peak hours. An AutoVots system without public transport would still remove 23% of the cars used today at peak hours. However, overall vehicle-kilometres travelled during peak periods would increase in comparison to today. For the TaxiBot with high-capacity public transport scenario, this increase is relatively low (9%). For the AutoVot car sharing without high capacity public transport scenario, the increase is significant (103%). While the former remains manageable, the latter would not be.

Reduced parking needs will free up significant public and private space

In all cases examined, self-driving fleets completely remove the need for on-street parking. This is a significant amount of space, equivalent to 210 football fields or nearly 20% of the kerb-to-kerb street space in our model city. Additionally, up to 80% of off-street parking could be removed, generating new opportunities for alternative uses of this valuable space.

Ride sharing with TaxiBots replaces more vehicles than car sharing with AutoVots

An AutoVot fleet requires more vehicles than a TaxiBot system to provide the same level of mobility. AutoVots also require considerably more repositioning travel to deliver that mobility.

The size of the self-driving fleet needed is influenced by the availability of public transport

Around 18% more TaxiBots and 26% more AutoVots are needed in scenarios without high-capacity public transport, compared to scenarios where shared self-driving vehicles are deployed alongside high-capacity public transport. Without public transport, 5 000 additional cars are required for the TaxiBot system and another 12 000 in the AutoVot system. Car-kilometres travelled would increase by 13% and 24% respectively.

Managing the transition will be challenging

If only 50% of car travel is carried out by shared self-driving vehicles and the remainder by traditional cars, total vehicle travel will increase between 30% and 90%. This holds true irrespective of the availability of

high-capacity public transport. Looking only at traffic during peak hours, the overall number of cars required increases in all but one scenario, namely TaxiBots with high-capacity public transport.

Policy insights

Self-driving vehicles could change public transport as we currently know it

For small and medium-sized cities it is conceivable that a shared fleet of self-driving vehicles could completely obviate the need for traditional public transport.

The potential impact of self-driving shared fleets on urban mobility is significant. It will be shaped by policy choices and deployment options

Transport policies can influence the type and size of the fleet, the mix between public transport and shared vehicles, and ultimately, the amount of car travel, congestion and emissions in the city.

Active management is needed to lock in the benefits of freed space

Shared vehicle fleets free up significant amounts of space in a city. Prior experience indicates that this space must be proactively managed in order to ensure these benefits are fully reaped. Management strategies can include restricting access to this space by allocating it to specified commercial or recreational uses, such as delivery bays, bicycle tracks or enlarged footpaths. Freed-up space in off-street parking could be used for urban logistics purposes, such as distribution centres.

Improvements in road safety are almost certain. Environmental benefits will depend on vehicle technology

The deployment of large-scale self-driving vehicle fleets will likely reduce both the number of crashes and crash severity, despite increases in overall levels of car travel. Environmental impacts remain tied to per-kilometre emissions and thus will be dependent on the adoption of more fuel-efficient and less polluting technologies. TaxiBots and AutoVots are in use 12 hours and travel nearly 200 kilometres per day, compared to 50 minutes and 30 kilometers for privately-owned cars today. More intense use means shorter vehicle lifecycles and thus quicker adoption of new, cleaner technologies across the car fleet.

New vehicle types and business models will be required

A drastic reduction in the number of cars needed would significantly impact car manufacturer business models. New services will develop under these conditions, but it is unclear who will manage them and how they will be monetised. The role of authorities, both regulatory and fiscal, will be important in guiding developments or potentially maintaining market barriers. Innovative maintenance programmes could be part of the monetisation package developed for these services.

Public transport, taxi operations and urban transport governance will have to adapt

Shared self-driving car fleets will directly compete with urban taxi and public transport services, as currently organised. Such fleets might effectively become a new form of low capacity, high quality public transport. This is likely to cause significant labour issues. Yet there is no reason why current public transport operators or taxi companies could not take an active role in delivering these services. Governance of transport services, including concession rules and arrangements, will have to adapt.

Mixing fleets of shared self-driving vehicles and privately-owned cars will not deliver the same benefits as a full TaxiBot/AutoVot fleet - but it still remains attractive

In all fleet-mixing scenarios, overall vehicle travel will be higher. Also, vehicle numbers will increase in three out of four peak hour scenarios. Improved traffic flow of automated cars could mitigate congestion up to a point. However, the public policy case for self-driving fleets alone (without high-capacity public transport) may be difficult to make based solely on space and congestion benefits, due to the increase in overall travel volumes. Nonetheless, even in mixed scenarios, shared self-driving fleets could be a cost-effective alternative to traditional forms of public transport, if the impacts of additional travel are mitigated. "All in" deployment of shared self-driving fleets may be easier in circumscribed areas such as business parks, campuses, islands, as well as in cities with low motorisation rates.

Introduction

For this study, we examined the potential outcomes of a radical change in urban mobility configuration that would result from the implementation of a shared and fully autonomous vehicle fleet.

To perform this assessment, we developed a new agent-based model to simulate the behaviour of all players of this system: First, the travellers, as potential users of the shared mobility system. Second, the cars, which are dynamically routed on the road network to pick-up and drop-off clients, or to move to, from, and between stations. Third, a dispatcher system tasked with efficiently assigning cars to clients while respecting the defined service quality standards, e.g. with regard to waiting time and detour time.

We based this analysis on a real urban context, the city of Lisbon, Portugal. We selected Lisbon as a case study due to the availability of data required to develop an agent-based simulation and because of its relative comparability with other European urban contexts.

This report is structured as follows: In chapter 1, we review similar research work to that which we carried out. In chapter 2 we characterise the Lisbon case study and highlight the main mobility-related attributes of that city for comparison with the outputs of our modelling exercise. Chapter 3 then briefly describes the simulation model and chapter 4 investigate several iterations of our basic shared-mobility scenarios. A discussion of policy implications from the results obtained concludes this study in chapter 5.

The work for this report was carried out in the context of a project initiated and funded by the International Transport Forum's Corporate Partnership Board (CPB). CPB projects are designed to enrich policy discussion with a business perspective. They are launched in areas where CPB member companies identify an emerging issue in transport policy or an innovation challenge to the transport system. Led by the ITF, work is carried out by in a collaborative fashion in working groups consisting of CPB member companies, external experts and ITF researchers.

The principal author of this report was Luis Martínez of the International Transport Forum who was also responsible for undertaking the modelling and analytical work upon which the report is based, some of which was completed during his time at the University of Lisbon. Special thanks to José Viegas who instigated and supervised this work. Substantial inputs were provided by Philippe Crist, who contributed to the project design and edited the final report, and to Maël Martinie who undertook valuable research in support of the work. Participating Corporate Partners in this report were Michelin and Nissan.

The project was coordinated by Philippe Crist and Sharon Masterson of the International Transport Forum.

1. Research review: Shared and self-driving car fleets

Cars are underused assets. They are mainly active during peak hours and rarely for more than 10% of the day – in fact, most are used for less than one hour a day. Much of their capacity is also underused since cars typically display low levels of occupancy in each trip – often with only one occupant. And despite this, they are highly valued assets – so highly valued that households put up with such levels of inefficiency in order to derive specific benefits relating to comfortable, door-to-door and schedule-less travel. Could this inefficiency be reduced while retaining these benefits?

Our work investigates the convergence of shared transport services, including car sharing/ride sharing and self-driving vehicle technology. The former has traditionally concerned largely informal and ad-hoc sharing (household car sharing, car pooling, etc.) but, starting in the 1980s, new models of co-operative-based and commercial car sharing emerged. These forms of car sharing allowed individuals to subscribe to shared fleets whose vehicles they reserve, access and use only when they need them. Pricing for these services is typically calculated on a per-hour or per-kilometre basis (or both). These services are situated somewhere between traditional car rental services and taxis and have proven popular in many urban areas since they allow individuals to have access to cars without necessarily owning one. With the arrival of ubiquitous internet access and dedicated app-based services, car sharing has quickly grown in popularity and sophistication and numerous successful services have been deployed around the world. At the same time, there has been an analogous development in terms of technological sophistication with ride-sharing services – especially for app-based on-demand services. These can take the form of taxi-like services or peer-to-peer real-time ride sharing. As with app-based car sharing, these forms of ride sharing have proven to be tremendously popular and pioneering companies in this field have generated billions of dollars in market capitalisation.

All of these services currently require a driver and so it seems interesting to examine what might be the next step in these services' evolution, namely, their integration with self-driving technology. This is not necessarily just a theoretical exercise – both Google and Uber have signalled both explicitly and implicitly that they see great potential for shared and autonomous vehicle fleets in both the car-sharing and ride-sharing modes. Several researchers have also examined the comprehensive impacts of the deployment of shared and self-driving vehicle fleets in various contexts. We focus on five of these in this section.

Mixing shared autonomous with traditional car fleets

A scenario developed by Fagnant and Kockelman (see box on p. 12) presents a model of a fleet of Shared Autonomous Vehicles (SAV) in a city of a size similar to that of Austin, Texas. The model has the following characteristics: each SAV travels autonomously, i.e. without human intervention, with at least one passenger to its final destination. In this model, there are no stops between origin and destination to board additional passengers, and no deviation occurs from the initial trip.

After each trip, the SAV moves on to the next traveller or repositions itself to a more favourable location for lower cost parking and faster future passenger service. This implies that there are no fixed stands that travellers have to reach to start their trips since the SAV comes to them. The fleet is comprised of traditional petrol-fuelled SAV sedans, i.e. no hybrid, electric or alternative-fuel vehicles were modelled. Finally, the authors consider only 3.5% of all trips as making use of the SAV network, the rest being made with conventional human-driven vehicles.

The modelling results suggest that each SAV would serve 31 to 41 persons per day, with an average waiting time below 20 seconds. Each SAV would replace nearly 12 conventional vehicles, and would lead to the elimination to 11 parking spaces per SAV in operation.

Overall distance travelled increased by 11% compared to a traditional human-driven self-owned fleet. This increase in travel distance was largely due to the relocation of the SAVs and the distance travelled to collect the next passenger. However, environmental impacts of the implementation of such a fleet are positive, with 5.6% less greenhouse gas emissions, 34% less carbon monoxide emitted, as well as a 49% reduction in volatile organic compound emissions, among others, compared to the traditional US light duty fleet. Emission reductions could be further reduced by considering a more intensive use of the SAV which would lead to a shorter life cycle for each vehicle (-1.5 to 2 years), hence an earlier replacement by more recent and less polluting vehicles. The use of an electric fleet could even further reduce emissions.

Among the identified limitations to this modelling exercise is the lack of a real-world context. Future modelling should be based on the geographical characteristics of a real urban area to provide more precise results capturing heterogeneous land use and travel patterns, seasonality and weekends. Other changes could include incorporating car-pooling options to improve the use of the SAVs as well as reducing overall distances travelled and associated environmental impacts. SAV impact on congestion is not measured in this paper.

Automated mobility-on-demand for Singapore

A 2014 study by Spieser et al. (see box on p. 12) explores the effect of a complete removal of the entire private vehicle fleet in Singapore, and its replacement by a shared self-driving fleet. The findings suggest that such a fleet could remove two thirds of the vehicles currently operating in Singapore while still delivering all of the trips currently made by private vehicles. The authors note several benefits of autonomous driving, such as better safety performance, an increase in the convenience and optimisation of trips, a decrease in congestion, lower overall costs, lower parking space requirements, etc.

While the case study focuses on shared self-driving vehicles, the authors note the findings could be extended to more general situations, such as shared vehicles with human drivers. However, the paper concludes that the most cost and time-effective option would be that of an automated mobility-on-demand (AMoD) system, as the shared self-driving model appears almost 50% cheaper than the model based on human-driven cars. However, such a system increases the overall distance travelled, as well as vehicle-use intensity, which may erode benefits linked to travel times and congestion.

Autonomous taxi system for New Jersey

Zachariah et al. (see box on p. 12) model the implementation of a fleet of autonomous taxis (ATaxis) in New Jersey, based on origin-destination trips derived from travel surveys. These trips approximate the real trips made by people in New Jersey every day. Passengers go to a station and take an ATaxi, which then brings them to the station nearest to their final destination. Other passengers can join the ride, provided that their destinations are located not too far from the destination of the first passenger.

Results suggest that there is significant ride-share potential. This potential is sensitive to relaxing the travel scheduling constraint away from the original trip. Average vehicle occupancy increases along with the increase of the waiting time at the station (to increase the chance that another passenger joins). It also increases when destinations of passengers are close to each other. The simulation shows that demand varies temporarily and spatially: The potential for ride sharing increases during peak hours, for example, and in locations such as railway stations. Taking that into account could make it possible for such a system to contribute to significant reductions of congestion in heavy-traffic areas, alongside a corresponding reduction in pollution.

Taxi pooling for New York City

This modelling work by Santi et al. at MIT's SENSEable City Lab (see box in p. 12) looks at the potential impact that the sharing of taxi rides could have on taxi fleet operation in New York City. It does so by looking at the detailed origin, destination and timing of every single taxi trip taken in the city over the course of a year and investigates which of these trips could have been shared, because riders were travelling from roughly the same areas to roughly the same destinations at approximately the same time.

A shared fleet is constructed in such a way that every real trip taken occurs in the model with no more than a five-minute delay to the real arrival time. Results suggest that the total number of kilometres driven by a taxi in New York City could be reduced by 40% with such a shared taxi system. This would consequently lead to large cuts in service costs, traffic congestion and emissions, as well as a reduction in fares paid by individual travellers. The authors conclude that it would be possible and efficient to implement a shareable taxi service in New York City.

The study also indicates that, even though the base case accounts for 150 million trips undertaken by 13 000 taxis in a large city, the model can be replicated in smaller cities up to a quarter of the size of the model city. The model does not take into account changes in the behaviour of passengers, who could respond to lower fares by increasing their use of the system. Also not fully addressed is the potential segmentation of the market, with a low end offering shared rides and a high end offering single passenger (or party) rides.

Transforming personal mobility: Three regional cases

This model exercise carried out by Burns et al at Columbia University (see box on p. 12) examines a shared, self-driving and centrally dispatched fleet of vehicles in three different environments: A mid-sized US city (Ann Arbor, Michigan), a low-density suburban development (Babcock Ranch, Florida) and a large and densely-populated urban context (Manhattan, New York). It uses travel survey-based data on average trip distances, trip-making rates (e.g. trips per hour) and travel speeds to help characterise travel in the regions studied. A combination of queuing, network and simulation models is used to calculate travel patterns and vehicle requirements. The modelling system generates trips to be serviced by a fleet of shared self-driving vehicles via a centralised dispatching systems that keeps track of the locations of all vehicles. The origins and destinations of trips are generated randomly over the whole of the region. Trips are requested at a constant average rate, times between requests are exponentially distributed and the single class of vehicle used in the model operates at a constant travel speed.

The study finds that for the 120 000 residents of Ann Arbor who travel less than 70 miles a day, the shared fleet could provide near instantaneous access to a vehicle servicing their request, but with only 15% of the vehicles currently needed to carry out these trips. However, overall travel would increase due to the need for repositioning vehicles. Similar findings emerge from the Babcock Ranch case study (3-4 000 vehicles for a projected population of 50 000 people). In the case of Manhattan, the study finds that a fleet of 9 000 taxis could replace all of the trips taken today by over 13 000 taxis with average waiting times of less than one minute, much lower than today.

The Columbia study also looked at costs from the consumer perspective in all three contexts. These increase with both travel and trip-making rates and differ depending on whether the vehicles used mimic traditional cars or are purpose-built for the shared and self-driving task. In the case of Ann Arbor, the authors estimate that the combined ownership, operating, parking and time value costs (linked to driving and searching for parking) are approximately USD 1.60 per mile for a conventional personally-owned car operated 10 000 miles per year. In contrast, they find that a shared and self-driving service using conventional cars would cost users about USD 0.41 per mile. Using a small, purpose-built 1-2 occupant car, reduces costs per trip mile even further, to USD 0.15 per trip mile. These are very significant cost

reductions. Shared self-driving fleets are found to cost more in the Babcock Ranch case study (USD 0.46 per trip mile) and in the Manhattan case study (USD 0.50 per trip mile) than in the Ann Arbor case study - but in both instances significantly less than the current usage of a conventional car fleet in those contexts.

This sample of recent publications emphasises the interest and relevance of research in shared and autonomous mobility services. It also provides an indication that findings in this research area are well-aligned, pointing to significant potential for reducing the number of vehicles and the parking space needed without reducing the level of mobility and service. Building on this work, the study presented in chapter 2 goes further by presenting detailed results of a full-scale simulation for a mid-size European city.

Studies reviewed in this section

L.D. Burns, W.C. Jordon, B.A. Scarborough (2013), "Transforming personal mobility", Earth Island Institute, Columbia University, 2013.

D. J. Fagnant, 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, March 2014.

P. Santi, G. Resta, M. Szell, S. Sobolevsky, S.H. Strogatz, and C. Ratti (2014), Quantifying the benefits of vehicle pooling with shareability networks, *Proceedings of the National Academy of Sciences of the United States of America (PNAS)*, 111(37), September, 2014.

K. Spieser, K. Ballantyne , K. Treleaven, R. Zhang, E. Frazzoli, D. Morton, M. Pavone (2014), *Toward a systemic approach to the design and evaluation of automated mobility-on-demand systems: A case study in Singapore. MIT Open Access Articles* (Forthcoming in Road Vehicle Automation, Springer Lecture Notes in Mobility series), 2014.

J. Zachariah, J. Gao, A. Kornhauser, T. Mufti (2013): "Uncongested mobility for all: A proposal for an area-wide autonomous taxi system in New Jersey", in: *Proceedings of Transportation Research Board Annual Meeting*, Washington D.C.

2. Case study: The city of Lisbon

In order to provide a frame of reference for our modelling exercise a detailed view of our test city's existing mobility patterns is necessary, along with its specific topology and other characteristics. This of course constrains the direct transferability of the results of the modelling exercise as they are firmly embedded in a unique local context. Nonetheless, insights derived from this exercise can provide an indication of the scale and direction of the expected impacts of the deployment of the proposed urban mobility system upgrade.

Our simulation exercise was based on data from the municipality of Lisbon, Portugal. Lisbon is the capital city of Portugal and is the largest city in the country with approximately 565 000 inhabitants in an area of 84.6 km². The city is the centre of the Lisbon Metropolitan Area (LMA), which has approximately 2.8 million inhabitants, representing roughly 25% of the Portuguese population, with an area of about 3 000 km^2 formed by 18 other municipalities.

The Lisbon Metropolitan Area generates over 5 million person-trips each day, of which 55% are commuting trips to go to work or study. Of this total activity in the LMA, about 1.2 million trips take place within the administrative boundaries of the Lisbon municipality. It is this set of trips that are examined in our model.

The inhabitants of the Lisbon municipality display a relatively low car ownership rate of 217 cars per 1 000 inhabitants and a low number of daily trips per inhabitant (1.9). Parking space is very scarce in the old heritage areas of the city, which deters car ownership. The low daily trip rate is related to the demographic profile of the city centre population - in some traditional city boroughs, more than 50% of the population is aged 65 years or older. Nevertheless, these figures align more or less with other similar European cities when controlling for wealth measured in GDP per capita.[1] Lisbon's main characteristics with regards to transport infrastructure and public transport services provision are outlined in Table 1.

The Portuguese capital has a well-established underground network that plays in important role in the daily mobility of inhabitants and workers of the metropolitan area, and especially within the municipality of Lisbon. The underground system was inaugurated in 1959 and has been growing ever since, currently covering a significant part of the Lisbon area, with some short links to suburban areas. The Lisbon Metro is a medium-sized network comprised of four lines covering about 43 kilometres of linear distance with 52 stations, transporting a total of 176.7 million passengers per year. Additionally, four main commuter rail lines connect the greater Lisbon area with the city centre. The relevance of these lines for the present study is limited, since only 13 stations are located within the city centre.

Recent data indicates that over 60 000 cars, 400 buses and 2 000 taxis circulate simultaneously during peak traffic periods in Lisbon, resulting in an average density of 60 vehicles per road kilometre (Martínez et al., 2014). This value is relatively high, especially when considering that a significant part of the road network is comprised of narrow streets.

---

[1] Based on analysis of UITP Millennium Cities Database; Mobility Observatory created by the Development Bank for Latin America.

Table 1. **Infrastructure and public transport provision in Lisbon**

Infrastructure provision		
Road	Road network density (km road/km^2)	12.58
	Motorway network density (km road/km^2)	0.24
Public transport	Underground network density (km line/km^2)	0.51
	Underground stations density (stations/km^2)	0.65
	Priority lanes for public transport (km road/km^2)	0.22
Parking	Street parking capacity	153 000
	Off-street parking capacity	50 000
	Parking availability (parking space/inhabitant)	0.38

Public transport services provision		
	Bus (thousands vehicle-km)	329.7
	Trains (thousands vehicle-km)	27.0
	Underground (thousands vehicle-km)	6 495.7

Source: Câmara Municipal de Lisboa, 2005.

Recent values estimated for the city indicate a maximum of 160 000 cars parked simultaneously, resulting in a very high utilisation rate of 78% of the available capacity (Câmara Municipal de Lisboa, 2005). This parking constraint has constrained car ownership rates in the heart of the city resulting in a more balanced mode share distribution in the city centre than in the whole of the metropolitan area (Table 2). While the latest census data reveals that 60% of all trips within the greater Lisbon metropolitan area are undertaken in a private car, this percentage drops to 40% in the city centre with 20% of trips there taking place by non-motorised means, principally by walking.

Table 2. **Share of transport modes for Lisbon**
(in %, 2011)

Modes	City of Lisbon	Lisbon Metropolitan Area
Private car	35.6	59.3
Motorcycle	2.2	1.2
Taxi	1.6	0.4
Bus	25.1	13.8
Walking/cycling	11.1	3.1
High-capacity public transport (underground or rail)	19.8	10.6
Car and high-capacity public transport	1.1	4.0
Bus and high-capacity public transport	3.5	7.7

3. Model description

In this work we developed an agent-based model that simulates the daily operation of a hypothetical shared mobility system in Lisbon. The model is based on real trip-taking activity, and the simulation takes place on Lisbon's real road network. The model is set up in such a way that the fleet of shared and self-driving vehicles deliver the same trips (in terms of origin, destination and timing) as those in a full-scale synthetic population generated on the basis of the Lisbon Travel Survey (Câmara Municipal de Lisboa, 2005), generalised to a grid of cells measuring 200 metres by 200 metres. A dispatcher system manages the centralised task of assigning mobility requests to cars using the location of shared self-driving vehicles, their current occupancy level and the location of clients as its main inputs. The model estimates trip routing on the basis of an algorithm that generates the lowest-cost path between any pair of nodes of the network.

Our model addresses the interaction between clients and vehicles, simulating their connection and how, in terms of timing and location, the services are performed. It does not include a dynamic traffic model which would simulate vehicle-level interactions amongst each other and with their environment. Our approach is based on a static representation of the traffic environment where origin-destination flows are allocated to a simple, topologically correct road network representation that accounts for per-link occupancy (and thus for speeds), by time of day.

Demand generation

On the basis of the Lisbon Travel Survey, we created a synthetic population of trips within the city, aggregated into the aforementioned grids. The synthetic travel simulation model we used was developed and calibrated for the LMA in previous studies, and its output contains all the trip origins and destinations spatially allocated at the census block level, as well as their timing for a synthetic weekday (Viegas and Martínez, 2010).

Each trip is characterised not only by its time of occurrence, origin and destination, but also by trip purpose and traveller's age. For all modes considered – shared vehicles (either ride sharing or car sharing), walking or cycling, and underground or suburban train – the trip is characterised by access time, waiting time, travel time and number of transfers between grids, where applicable.

We adopted a rule-based approach in order to specify a simplified and restricted mode choice process. All short trips under one kilometre were considered to be undertaken by foot or by bicycle (in reality, almost all of these would be taken by foot). Other trips were either assigned to the underground or to the shared and self-driving mode alternative. Trips where both the origin and destination were in proximity to an underground station, and where the total trip required no more than one transfer within the underground network, were assigned to that mode. The remaining trips not well served by underground connections or by walking were then assigned to the shared and self-driving alternative considered for a variety of scenarios. In those scenarios where we considered that no underground services were available, all trips of more than one kilometre in length were assigned to the shared and self-driving mobility option.

When the model chooses the shared and self-driving option, a new user (agent) is generated in the simulation environment, with a departure node, an arrival node and a starting time. Currently, one user is equivalent to one trip, i.e. users do not cluster in parties at the outset of their trip, though they do share vehicles once the ride-sharing simulations are underway.
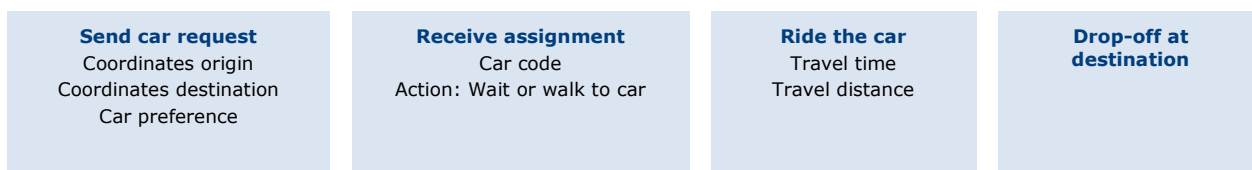
Trip generation by users

In the simulation environment, a trip is generated when a user requests a departure from a point towards another point. The model accounts for the simulation parameters (ride sharing or car sharing) and accounts for waiting time, detour time and arrival time tolerances that are defined for the model run. The dispatcher then finds, in real time, the best possible routing and assigns one of several available vehicle types to carry out the trip in either a car-sharing or ride-sharing mode.

The user then waits for the car or accesses a specified pick-up location and boards the vehicle. When the vehicle arrives at its destination, the user exits the system and a set of indicators are generated in a trip log so that they can be used for ex-post system evaluation.

Figure 1. **Flowchart of the agent-based model for trip generation**

Send car request Coordinates origin Coordinates destination Car preference	**Receive assignment** Car code Action: Wait or walk to car	**Ride the car** Travel time Travel distance	**Drop-off at destination**

Car configurations

Idle cars are located in 60 stations around the city. Whenever a car is empty and not immediately dispatched to a new trip, it relocates itself to a station (in the TaxiBot ride-sharing system) or parks itself (in the AutoVot car-sharing system). Active cars follow the shortest path and minimise travel time for their route assignment, taking into account hourly link-based road speeds.

Our approach assumes a fleet comprised of three types of vehicles, based on available passenger capacity: two-passenger, five-passenger and eight-passenger cars. We did not specify the type of propulsion technology - this is of minor relevance for the present study, since emissions or energy use were not modelled. We did however test the sensitivity of the system to added recharging time and fleet requirements for battery electric vehicles.

Role of the mobility dispatcher

The mobility dispatcher is an entity that defines a set of rules for matching cars to users, centralising all real-time information required to produce and monitor these trips. The choice of which car to match with a user request takes into account a time-minimisation principle that applies not just to the requesting user but also to those already underway.

We defined a number of parametric constraints on the runs that must be satisfied for each route proposed by the mobility dispatcher. These include constraints such as the requirement that the generated trips commence with no more than a five-minute delay on the standard "base case" trip derived from the original trip data (Viegas and Martínez, 2010). Other constraints relate to the maximum number of users that can share a vehicle (eight people maximum) and a maximum allowable increase in time (no more than an additional 20% travel time compared to the original trip, capped at a maximum of ten minutes) and distance (maximum of 20% of current travel distance, capped at 2 kilometres) when compared with the direct connection. This means for instance that in the worst case, a trip that took ten minutes in the base case would take seven minutes more in the simulation – a five-minute waiting period for the trip to commence plus an additional two-minute travel time (20% of the original ten minute travel time).

In reality, simulated travel times for the shared and self-driving scenarios are often lower, especially for trips formerly taken by bus, and are rarely significantly longer than trips in the base case. Furthermore, travel times in the base-case scenario do not account for time required to access vehicles and time spent searching for parking for car users. Both would lengthen car travel times for the base case and reduce the gap with travel times in the simulated shared and self-driving system.

The simulation has a dynamic graphical display to visualise the model workflow during the simulation (Figure 2). This component allows inspecting the dynamic variables of each client and car during the simulation.

Figure 2. **Visualisation of shared self-driving car simulation for Lisbon**

4. Testing shared-mobility scenarios

We devised several scenarios for the deployment of shared and self-driving fleets in Lisbon. These scenarios were generated by varying four principal parameters in our simulation:

- The **mode** of shared and self-driving operation for the simulated fleet, namely either car sharing or ride sharing.

- The availability of **high-capacity public transport** or not.

- The **penetration rate** of the shared and self-driving fleet – either 100% or only 50% of all trips.

- The **time period** considered for the simulation – either all day on a weekday or only during peak travel periods.

We investigated two system configurations for this fleet, both of which consist of fully self-driving vehicles:

- A **ride sharing system**, where travellers share time and space resources by travelling in the same car simultaneously up to the capacity limit of the vehicle. The cars may either be privately owned by one rider or from a car fleet company. We labelled this a "Taxi Robot" system, or **TaxiBot.**

- A **car sharing system**, where travellers share time resources by travelling in the same car sequentially. In this case, car fleets are normally owner by a car fleet manager, although there are also some incipient peer-to-peer experiments. We labelled this an "Automated Vehicle Robot" system, or **AutoVot**.

The decision to simulate the absence of a high-capacity public transport option was motivated by the need to see how well the shared and self-driving option could absorb these trips. In the present case, the high-capacity public transport option was an underground but other high-capacity public transport solutions such as commuter rail, Bus Rapid Transit (BRT) and Light Rail Transit (LRT) could also be used if they present a similar level of station density as Lisbon (0.65 stations/km^2).

We also modelled scenarios that account for only a 50% penetration rate for the shared and self-driving fleets. In these scenarios, we randomly assigned 50% of the trips not undertaken by high-capacity public transport or walking to traditionally operated conventional cars and 50% to the shared and self-driving fleet. All of these scenarios were run under the generic waiting time, detour time and distance and vehicle capacity constraints outlined above.

A single simulation run of a regular day was performed to analyse the mobility outcomes for the different scenarios. Although these results may present some variability between trials, the stability for the main extracted indicators is high.

Modal shares in different scenarios

As the modal-choice model is encompassed in a rule-based decision model, the obtained shares for the different modes are very stable for all tested scenarios. The results suggest an increase of high-capacity public transport use when compared with the current pattern of travel in Lisbon. This deviation results from current bus users being assigned to the underground, even though they may prefer the former to the latter in reality. The one-kilometre threshold we set for assigning all trips to walking is low, as the real share for walking in Lisbon is considerably higher than for comparable cities.

In the scenarios where all motorised trips were previously carried out by non-high-capacity public transport, the TaxiBot/AutoVot fleets reach a share of 70% (if complemented by high-capacity public transport) and 92% (in absence of the latter), as shown in Table 3.

Table 3. **Mode share distribution for different TaxiBot and AutoVot scenarios**
(based on Lisbon, 24-hour average, weekday)

			Mode share (%)		
		Transport mode	Public Transport	Walking and cycling	Cars
		Baseline	15	18	48*
100% shared self-driving fleet	Ride sharing (TaxiBot)	No high-capacity public transport	0	8	92
		With high-capacity public transport	22	8	70
	Car sharing (AutoVot)	No high-capacity public transport	0	8	92
		With high-capacity public transport	22	8	70
50% private car use for motorised trips	Ride sharing (TaxiBot)	No high-capacity public transport	0	8	46
		With high-capacity public transport	22	8	35
	Car sharing (AutoVot)	No high-capacity public transport	0	8	46
		With high-capacity public transport	22	8	35

* Baseline assessed from share of private cars

Impact on fleet size

The impacts on fleet size on the different scenarios are set out in Table 4. It shows that shared self-driving fleets have the potential to drastically reduce the number of vehicles necessary to deliver the same travel as today's fleet. Under a ride-sharing TaxiBot configuration supported by high-capacity public transport and modelled over a 24-hour weekday, 90% of vehicles could be removed from the streets while still delivering nearly the same level of mobility as before in terms of travel origins, destinations and length of trip.

Even more striking is that only approximately 5 000 shared self-driving TaxiBots would be required to handle all of the trips currently modelled for the Lisbon Metro in our exercise. A car-sharing AutoVot operation displays less potential for fleet reduction than a ride-sharing TaxiBot system, since the former requires more vehicles and much more repositioning travel to deliver the same level of service. If there is no high-capacity public transport option such as an underground system available, the 50% penetration scenarios for TaxiBots or AutoVots result in more vehicles being required than in the base case. Even in the presence of an underground, there is only a limited reduction of the number of private cars. This suggests that transition scenarios may not lead to expected levels of car fleet reduction in the short term.

Table 4. **Fleet size for different TaxiBot and AutoVot scenarios**
(% of current Lisbon car fleet, 24-hour weekday average)

			Fleet size	% of baseline
		Baseline	203 000	
100% shared self-driving fleet	Ride sharing (TaxiBot)	No high-capacity public transport	25 917	12.8
		With high-capacity public transport	21 120	10.4
	Car sharing (AutoVot)	No high-capacity public transport	46 249	22.8
		With high-capacity public transport	34 082	16.8
50% private car use for motorised trips	Ride sharing (TaxiBot)	No high-capacity public transport	13 265 + 194 537*	102.4
		With high-capacity Public transport	10 900 + 147 767*	78.2
	Car sharing (AutoVot)	No high-capacity Public transport	22 887 +194 275*	107.0
		With high-capacity Public transport	18 358 +148 050*	82.0

* = shared + private cars

Impact on travel volume

Shared and self-driving fleets hold much promise for reducing the number of cars in our cities. The same cannot be said for reductions in car travel, however. Table 5 looks at the impacts on cumulative travel volumes under the various shared and self-driving fleet scenarios over the duration of a weekday.

Total distance travelled increases under all of the scenarios we studied. Overall car kilometres travelled over the course of the day rose 6% in the TaxiBot system with high-capacity public transport and nearly doubled (+89%) in the AutoVot scenario with no public transport.

This increase can be explained by several factors. In all scenarios, we assume that shared self-driving fleets completely replace current bus travel in the city. In the base case, average bus occupancy over the course of the day is low (20%), and it is plausible that these passengers could be better served by a fleet of shared self-driving cars. The diversion of bus passengers accounts for approximately 30% of the final car-kilometres travelled in the ride-sharing scenarios and nearly 50% of the car-kilometres travelled in the car-sharing scenarios. The remaining travel increment is due to the repositioning of empty cars in all scenarios, as well as to detours for passenger pick-ups and drop-offs in the ride-sharing scenarios.

The mixed-fleet scenarios with 50% shared self-driving vehicles and 50% conventional cars result in a very significant increase in car travel – up to 90% in the case of an AutoVot system without public transport. These results reiterate the challenges for operating a shared-mobility system in a hybrid configuration.

Table 5. **Weekday travel volumes under different TaxiBot and AutoVot scenarios**
(for 24-hour weekday)

			Car-kilometers (millions)	% of baseline
		Baseline	3.8	
100% shared self-driving fleet	**Ride sharing (TaxiBot)**	No high-capacity public transport	4.62	122.4
		With high-capacity public transport	4.01	106.4
	Car sharing (AutoVot)	No high-capacity public transport	7.15	189.4
		With high-capacity public transport	5.44	144.3
50% private car use for motorised trips	**Ride sharing (TaxiBot)**	No high-capacity public transport	6.04	160.2
		With high-capacity public transport	4.90	129.8
	Car sharing (AutoVot)	No high-capacity public transport	7.20	190.9
		With high-capacity public transport	5.69	150.9

Similarly to all-day travel, car-kilometres travelled at peak hours increase compared to the base case in all scenarios considered (Table 6). For the TaxiBot scenario with high-capacity public transport, this increase is relatively low (9%). For the AutoVot scenario without high capacity public transport, the increase is significant (103%). While the former increase remains manageable, the latter would definitely not be.

Figure 3 shows the distribution of car-kilometres travelled over a 24-hour weekday for selected scenarios. It shows that peak-hour travel measured in car-kilometres increases in the TaxiBot and AutoVot scenarios. It also shows that the morning peak shifts from a maximum around 8 a.m. in the base case to a maximum around 9 a.m. in shared self-driving scenarios. The afternoon peak displays a similar but weaker shift.

Table 6. **Peak-hour travel volumes under different TaxiBot and AutoVot scenarios**
(weekday morning peak, 7 a.m. to 10 a.m.)

			Car-kilometers (millions)	% of baseline
		Baseline	1.04	
100% shared self-driving fleet	**Ride sharing (TaxiBot)**	No high-capacity public transport	1.30	125.3
		With high-capacity public transport	1.13	108.8
	Car sharing (AutoVot)	No high-capacity public transport	2.11	203.2
		With high-capacity public transport	1.60	154.6
50% private car use for motorised trips	**Ride sharing (TaxiBot)**	No high-capacity public transport	1.67	167.5
		With high-capacity public transport	1.36	135.8
	Car sharing (AutoVot)	No high-capacity public transport	2.04	197.0
		With high-capacity public transport	1.62	155.7

The increase in peak-hour car-kilometres travelled is neither uniform throughout the city nor on all road classes. We investigated how car-kilometres travelled are allocated across the city and among road classes by time of day. This was calculated for both the base case and some of the TaxiBot and AutoVot scenarios. The base case assessment we undertook covers the movement of private cars, motorcycles and taxis. It does not account for bus travel, since the spatial allocation of vehicle-kilometres was not readily available. This may have an impact, since bus travel represents 13% of total vehicle-kilometres travelled in Lisbon.

Figure 3. **Time distribution of travel volumes for selected TaxiBot and AutoVot scenarios**
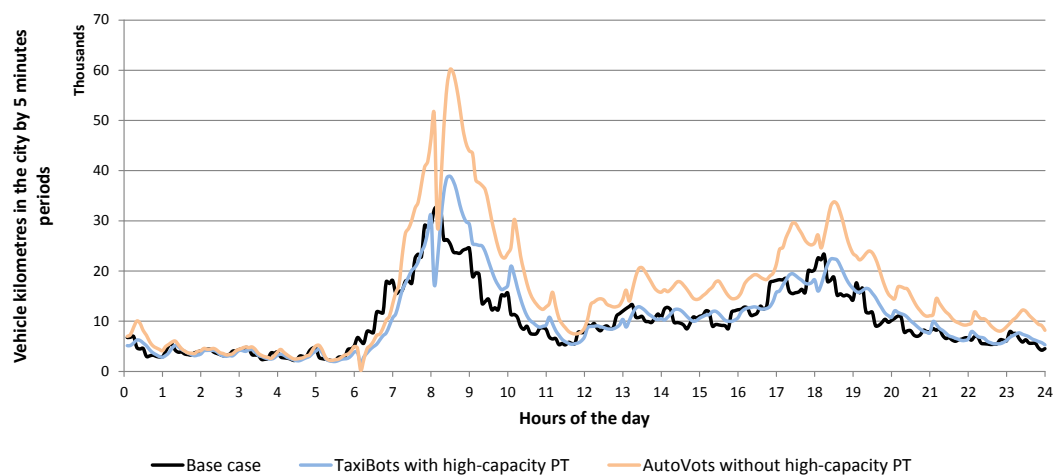(for 24-hour weekday)

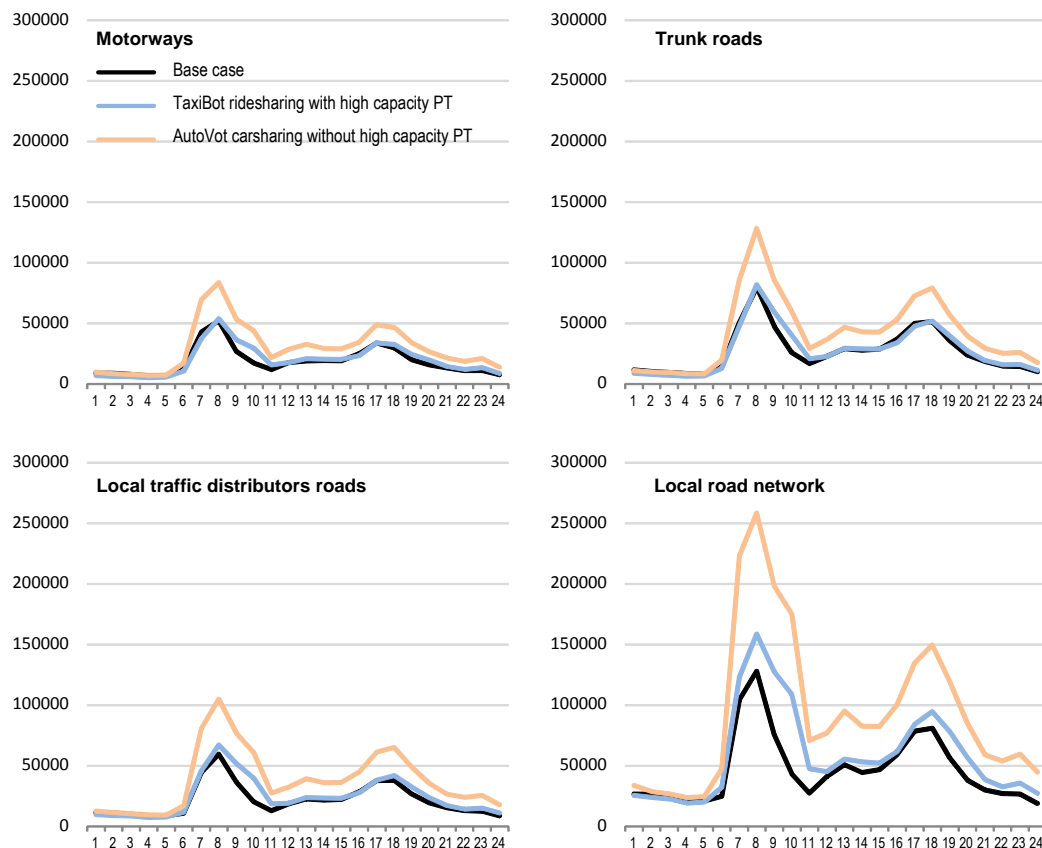Figure 4. **Time distribution of travel volumes by road class for selected TaxiBot and AutoVot scenarios**

Figure 4 compares the distribution of car-kilometres travelled over the course of a typical weekday for various scenarios by road class. It shows that the modest rise in travel volume in the TaxiBot scenario with high-capacity public transport does not contribute significantly to peak hour travel on all but the local road network. The increase in activity on local road networks is due to pick-up and drop-off movements, which seem more frequent during the morning peak. The figure also shows that there is some peak travel spreading on all road networks, which is a result of vehicle relocation to stations after peak demand periods. This is a relevant result from an operational perspective, as it implies that strategies will be needed to handle relocation following periods of intense use in anticipation of the next peak. In comparison to the TaxiBot scenario, the AutoVot scenario without high-capacity public transport shows an increase in travel volumes as well as a combined spreading and shift of the peak period on all road networks.

The fact that both the TaxiBot and AutoVot scenarios lead to an increase in travel on local arterials and the local road network may imply changes in the performance and characteristics of those networks. In terms of performance, current road occupancy at peak hours (in terms of percentage of technically feasible road capacity used) is generally below 40%, and less than half of that for the local road network (Table 7). In the TaxiBot plus high-capacity public transport scenario, peak-hour road occupancy barely changes for all but local roads, and even for the latter the increase is relatively modest at 23%.

This is not the case for the AutoVot scenario without high-capacity public transport. In this scenario, road occupancy increases by 40 to 50% for all road classes, with the strongest growth occurring on local road networks. This implies poorer performance and possibly congestion. It may also mean a change in the

nature of local street traffic resulting from the presence of additional traffic on what were otherwise quieter, less-used roads. Such an increase in traffic could have a negative impact on the attractiveness and livability of local streets, as it reduces their availability for non-transport use.

Table 7. **Road occupancy by road class during morning peak for selected TaxiBot and AutoVot scenarios**
(8 a.m. weekday, in % of technical capacity used)

Road hierarchy	Base case	TaxiBots with high-capacity public transport	AutoVots without high-capacity public transport
Highways	35.8	36.9	58.6
Trunk roads	42.5	44.5	68.9
Local traffic distributor roads	34.3	38.7	60.3
Local road network	16.7	21.9	36.0

Figure 5 depicts the spatial distribution of trips during the morning peak for the TaxiBot plus high-capacity public transport scenario in comparison to the base case. It shows that some roads and certain areas of the city may even see a drop in traffic during peak hours in the TaxiBot scenario, while others will experience a slight increase. The visualisation reinforces our two main findings regarding the traffic impacts of the TaxiBot scenario: It shows a marginal negative impact in congestion at certain bottlenecks, but also that overall traffic fluidity is largely preserved. It equally reveals the increase in traffic on local networks where traffic was largely absent previously, and thus the potential increase of traffic conflicts with walking and cycling in these areas. It should also be noted that the areas registering the largest increases in traffic volumes are those least well connected to high-capacity public transport. This underscores the complementarity of these two systems in our scenarios.

Figure 5. **Spatial distribution of the variation of peak hour travel volumes for TaxiBot system in Lisbon**
(weekday 8-9 a.m., TaxiBot plus high-capacity public transport scenario, vehicle-kilometres)
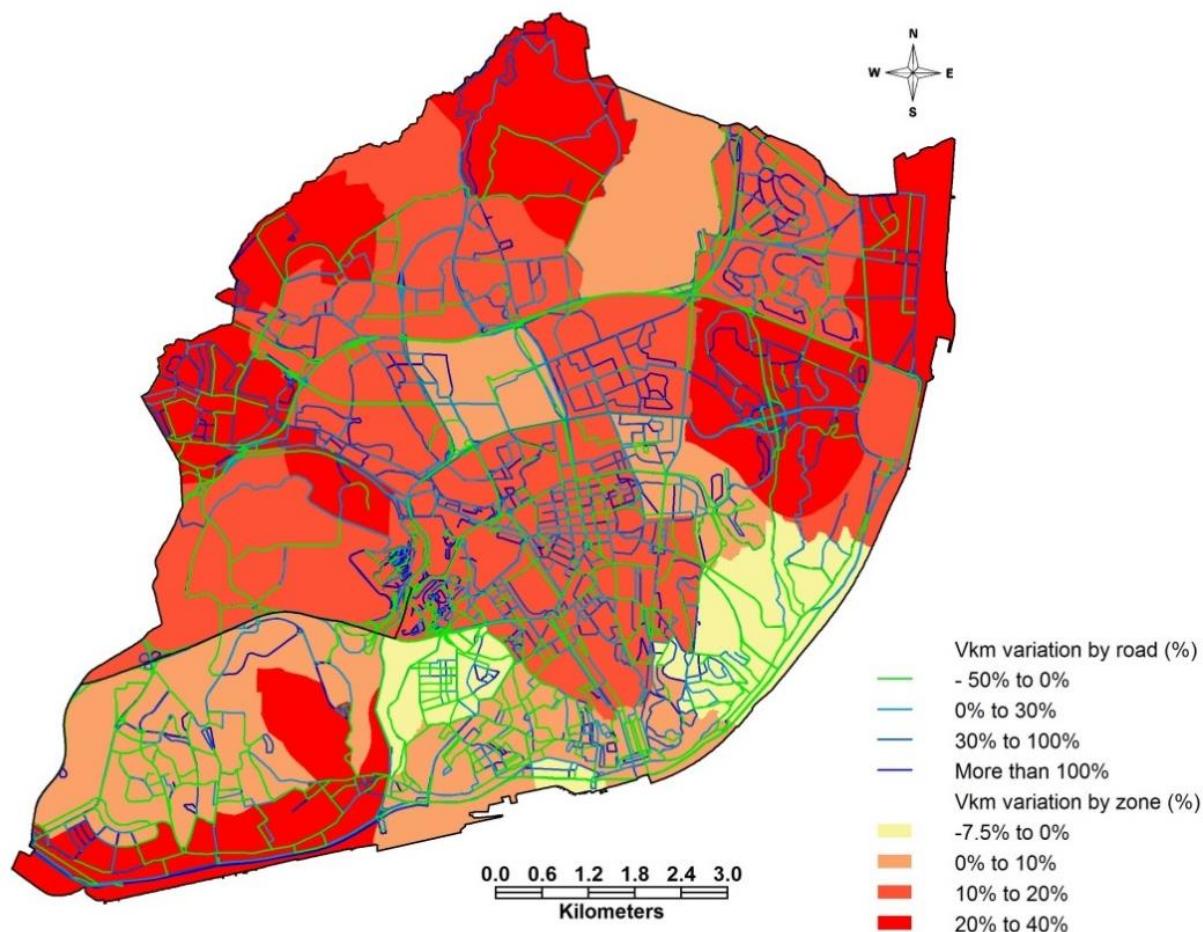
Vehicle fleet requirements at peak hours

In order to assess how the TaxiBot and AutoVot fleets performed at peak travel periods, we carried out an aggregate assessment of the number of vehicles circulating in the city during the morning peak. This part of the analysis does not factor in capacity limitations at link levels, but allows us to gauge peak-hour fleet requirements and travel volumes, and to compare these with the baseline (see Table 8).

Even at peak hours, the car reduction effect of shared self-driving fleets is important. We found that significantly fewer cars than today would be travelling at peak hours, while service levels are largely maintained compared to the base case. In the TaxiBot plus public transport scenario, 65% fewer cars are circulating at peak hours. Even in the least car-reducing scenario – AutoVots without public transport – 23% fewer cars were required.

Testing of the mixed-fleet scenarios indicates that ride-sharing systems may lead to reductions in peak loads when combined with high-capacity public transport. Under these scenarios, increase of relocation operations of car-sharing systems may produce a significant increase of traffic that may be absorbed with difficulty in already very congested urban areas.

Table 8. **Number of cars travelling during morning peak in selected TaxiBot and AutoVot scenarios**
(weekday morning, Lisbon)

			Cars travelling at peak hours	% of current max. flow
		Baseline	60 000	
100% shared self-driving fleet	**Ride sharing (TaxiBot)**	No high-capacity public transport	25 867	43.1
		With high-capacity public transport	21 105	35.2
	Car sharing (AutoVot)	No high-capacity public transport	46 011	76.7
		With high-capacity public transport	33 975	56.6
50% private car use for motorised trips	**Ride sharing (TaxiBot)**	No high-capacity public transport	13 173 + 57 499*	117.8
		With high-capacity public transport	10 890 + 43 675*	90.9
	Car sharing (AutoVot)	No high-capacity public transport	22 768 + 57 421*	133.6
		With high-capacity public transport	18 305 + 43 759*	103.4

* = shared + private cars

These results suggest that shared self-driving fleets may face some transition issues due to the presence of legacy fleets. Unless increases in travel at peak hours are managed, it may be difficult to make a public policy case for self-driving fleets based solely on freed space and congestion benefits in the presence of legacy conventional cars. Nonetheless, even in these scenarios, these fleets could potentially represent a cost-effective alternative to public transport if the impacts of additional travel are mitigated.

Another relevant output of the scenarios we tested is the peak travelling ratio, assessed as the ratio between the number of cars travelling at 8 a.m. and at 3 p.m. (Table 9). The values for this ratio indicate that car sharing under an AutoVot system leads to a greater number of cars on the streets at peak hours compared to ride sharing with a TaxiBot system. The peak travelling ratio is exacerbated when considering mixed-fleet scenarios, suggesting greater concentration of vehicles during peak periods and consequently greater congestion levels for respective peak travel periods.

Table 9. **Ratio of cars travelling in-peak and off-peak in selected TaxiBot and AutoVot scenarios**
(Weekdays, 8 a.m. to 3 p.m.)

			Peak/Off-peak factor
		Baseline	3.30
100% shared self-driving fleet	**Ride sharing (TaxiBot)**	No high-capacity public transport	2.65
		With high-capacity public transport	2.69
	Cars haring (AutoVot)	No high-capacity public transport	3.01
		With high-capacity public transport	3.17
50% private car use for motorised trips	**Ride sharing (TaxiBot)**	No high-capacity public transport	2.56
		With high-capacity public transport	2.53
	Car sharing (AutoVot)	No high-capacity public transport	3.04
		With high-capacity public transport	3.16

Impact on parking and street space

Vehicles use a considerable amount of space in cities, both when moving and when parked. Accordingly, we examined the impacts that shared self-driving vehicles would have on parking space requirements. Table 10

presents the maximum parking requirements implied by each scenario. It shows that there is an extremely large potential for reducing both on-street and off-street parking spaces, and that this is the case for all shared and self-driving fleet scenarios.

Considering that off-street parking represents 50 000 spots in the baseline case and that the most parking-intensive scenario (car sharing without public transport) would require 25 621 spots, on-street parking spots could be totally removed from the streets, whatever the scenario considered. This would allow the reallocation of 1 530 000 m² to other public uses[2], equivalent to almost 20% of the kerb-to-kerb street area in Lisbon or 210 football fields. This freed-up space could be dedicated to non-motorised transport modes, delivery bays, parklets or other recreational and commercial uses.

In the most favourable scenario for parking, a maximum number of 8 901 parking spots would be necessary. Taking 10 000 as a proxy of the total number of required parking spaces, 40 000 off-street parking spaces could then be re-allocated, i.e. 1 200 000 m² or the equivalent of almost 170 football fields. This space could be used in innovative ways, e.g. as centrally located urban logistics distribution centres.

Table 10. **Maximum number of parked vehicles for different TaxiBot and AutoVot scenarios**
(for 24-hour weekday)

			Max. Parking requirements	% of baseline
		Baseline	160 000	
100% shared self-driving fleet	**Ride sharing (TaxiBot)**	No high-capacity public transport	11 563	7.2
		With high-capacity public transport	8 901	5.6
	Car sharing	No high-capacity public transport	25 621	16.0
		With high-capacity public transport	17 110	10.7
50% private car use for motorised trips	**Ride sharing (TaxiBot)**	No high-capacity public transport	5 928 + 153 122*	99.4
		With high-capacity public transport	4 622 + 116 689*	75.8
	Car sharing	No high-capacity public transport	12 705 + 153 330*	103.8
		With high-capacity public transport	9 561 + 116 467*	78.8

* = shared + private cars

However, when testing the scenarios with mixed fleets of TaxiBots/AutoVots and traditional, privately-owned cars, the results are less promising. The reductions in parking space requirements are small in the TaxiBot system (up to 25%) and even negative in the AutoVot system. These findings suggest that shared and self-driving fleets operating in parallel with private conventional car fleets may lead to even higher parking requirements than today in the absence of bus services.

Impact on vehicle use

Another relevant aspect of shared and self-driving systems is their potential impact on more efficient vehicle operation – and in particular their potential to reduce unproductive down-time for vehicles. Currently, cars in the city of Lisbon are used for approximately 50 minutes per day, which means that they lie unused for 95% of the day.

---

[2] Assuming 10m² for each on-street parking space and 30m² for each off-street parking space (including lanes, ramps, etc.).

Table 11 displays how much time vehicles are not being used productively.[3] In all scenarios we see a decrease in the amount of unproductive idle time for vehicles. Unproductive vehicle downtime drops from 95% to less than 40% in all scenarios, and down to 27% in some scenarios.

Table 11. **Shares of idle time for different TaxiBot and AutoVot scenarios**

(cars parked, driving to depot or to initial client; for 24-hour weekday)

			Idle time (in %)
		Baseline	96
100% shared self-driving fleet	**Ride sharing (TaxiBot)**	No high-capacity public transport	27.1
		With high-capacity public transport	27.2
	Car sharing (AutoVot)	No high-capacity public transport	39.3
		With high-capacity public transport	35.2
50% private car use in motorised trips	**Ride sharing (TaxiBot)**	No high-capacity public transport	22.8
		With high-capacity public transport	23.6
	Car sharing (AutoVot)	No high-capacity public transport	39.4
		With high-capacity public transport	38.6

Distribution of vehicle types

One of the outcomes that we tracked was the distribution of vehicle types used to carry out the shared and self-driving trips (see Table 12). The results suggest that the shared self-driving fleet composition will be dominated by 3-5 passenger vehicles followed by small vehicles carrying one to two passengers.

Table 12. **Share of travel by vehicle type for different TaxiBot and AutoVot scenarios**

(24 hours, weekdays)

			Share of total car-kilometres (%) for:		
			Cars for 1-2 passengers	Cars for 3-5 passenger	Cars for 5-8 passengers
100% shared self-driving fleet	**Ride sharing (TaxiBot)**	No high-capacity public transport	34.4	60.3	5.3
		With high-capacity public transport	35.2	58.7	6.1
	Car sharing (AutoVot)	No high-capacity public transport	100	0.0	0.0
		With high-capacity public transport	100	0.0	0.0
50% private car use in motorised trips	**Ride sharing (TaxiBot)**	No high-capacity public transport	31.8	48.5	19.7
		With high-capacity public transport	33.6	47.1	19.3
	Car sharing (AutoVot)	No high-capacity public transport	100	0.0	0.0
		With high-capacity public transport	100	0.0	0.0

Figures for car sales in Lisbon in 2009 (Figure 6) indicate that many city dwellers own small cars. This suggests that the expected makes and models present in the Portuguese market, in terms of vehicle size, would not change drastically under shared and self-driving scenarios.

---

[3] In the sense that they are not providing work, e.g. that they are parked, are driving to a depot or to an initial client. We do not consider an empty vehicle repositioning itself to pick up a client as unproductive or idle.

Figure 6. **Sales of new cars in Lisbon by car type**
(2009)
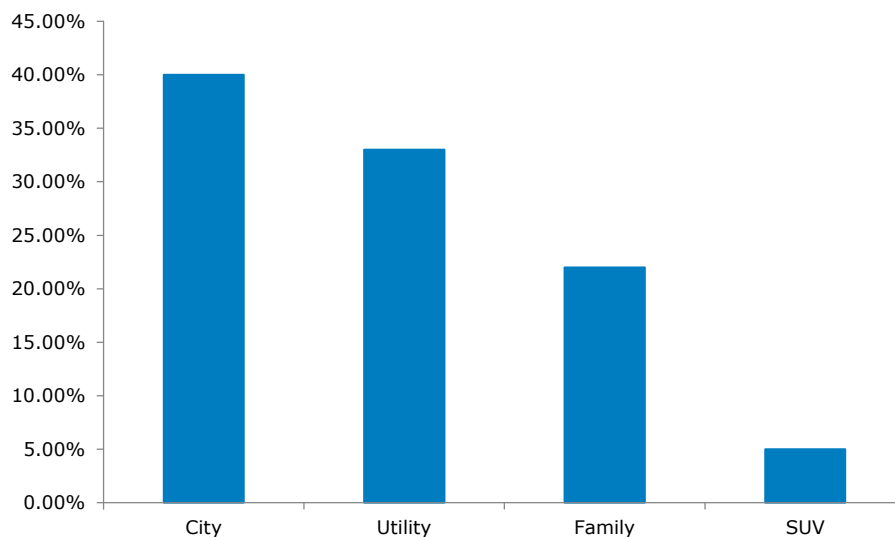
Source: Portuguese Car Trade Association (ACAP), 2009.

Impact of a fully electric vehicle fleet

We noted earlier that shared and self-driving fleets would likely lead to an increase in overall volumes of travel. Emissions of greenhouse gases and pollutants generated by this travel could be mitigated through the deployment of low-emission and energy-efficient drivetrain technologies. In the extreme case, fully electric fleets would eliminate tank-to-wheel emissions from all of the car travel in our scenarios. With current technologies, however, full battery-electric vehicles have limited range compared to conventional fossil-fuelled cars and would require additional off-line time for recharging.

We modelled additional vehicle requirements required by shared self-driving fleets consisting of electric vehicles to gauge the impact of re-charging times and reduced travel range. We assumed a fast battery recharging time of 30 minutes and vehicle autonomy of 175 kilometres. We found that the impact on fleet size of the deployment of a shared self-driving fleet of fully electric vehicles was minimal (+2%).

Changes in waiting and travel times

We assessed the average waiting and travel time resulting from shared mobility services via TaxiBots and AutoVots, and compared these with the baseline scenario for public transport and private car use. We found that this resulted in a significant reduction in average waiting and travel times (Table 13). These reductions are the result of the more personalised door-to-door services offered by TaxiBots and AutoVots, notably for those trips previously taken by bus, and improved travel times, especially in peak periods.

Relaxing our constraint on maximum waiting time to ten minutes (from five minutes in the initial model) leads to an 11% reduction of the required car fleet. Travel volume, measured in car-kilometres, would drop by a mere 3% in the best scenarios, since part of the additional time available would be used to pick-up and drop-off passengers in further-away locations. However, relaxing the waiting-time constraint would also result in a 27% increase of average party size in ride-sharing systems. It is therefore conceivable that such

a configuration would lead to a price reduction for clients if the service price for each user were linked to the number of riders sharing a vehicle.

Table 13. **Average waiting and travel time for different TaxiBot and AutoVot scenarios**
(variation to current public transport waiting time and private car travelling time in Lisbon)

			Urban density Lisbon density (x1) (7 400 inhab./km^2)			
			Waiting time (min)	Change (%)	Travel time (min)	Change (%)
		Baseline	26.37		18.30	
100% shared self-driving fleet	**Ride sharing (TaxiBot)**	No high-capacity public transport	3.73	-85.9	15.11	-17.4
		With high-capacity public transport	3.79	-85.6	15.93	-13.0
	Car sharing (AutoVot)	No high-capacity public transport	3.60	-86.4	11.36	-37.9
		With high-capacity public transport	3.03	-88.5	11.43	-37.5
50% private car use for motorised trips	**Ride sharing (TaxiBot)**	No high-capacity public transport	3.85	-85.4	18.08	-1.2
		With high-capacity public transport	4.05	-84.6	19.63	7.3
	Car sharing (AutoVot)	No high-capacity public transport	3.31	-87.5	11.46	-37.4
		With high-capacity public transport	3.38	-87.2	12.61	-31.1

The statistical distribution of waiting times across all clients obtained from the dispatching algorithm for a maximum waiting time of five minutes is presented in Figure 7. It shows that the most frequent waiting time is 3.3 minutes and the average 3.7 minutes for the TaxiBot scenario with high-capacity public transport. The shape of the distribution reveals two different branches of the dispatching algorithm: one branch associated with the assignment of initial passengers to an empty TaxiBot with a peak at 3.3 minutes, and another branch associated with the optimal allocation of a partially occupied vehicle to a client with a peak at 4.2 minutes.

Figure 7. **Distribution of waiting times for TaxiBot scenario with high-capacity public transport**
(curve adjusted and smoothed)

Waiting time [min]

In the base case, door-to-door travel times for bus and underground users include time to access the stop or station as well as some waiting time to reflect schedules. For car travel, no waiting or access time was assumed, even though parking may not be available in proximity of the point of departure or arrival. Search time for parking spaces was also not included. In this respect, our base case underestimates car travel times. Because of these assumptions, it is reasonable to expect that users of public transport will see an improvement in their door-to-door travel time in TaxiBot/AutoVot scenarios, while some car drivers may experience an increase in door-to-door travel times. The overall impacts on travel times would therefore be dependent on the share of car drivers relative to public transport users.

Figure 8. **Distribution of daily travel time under different TaxiBot and AutoVot scenarios**
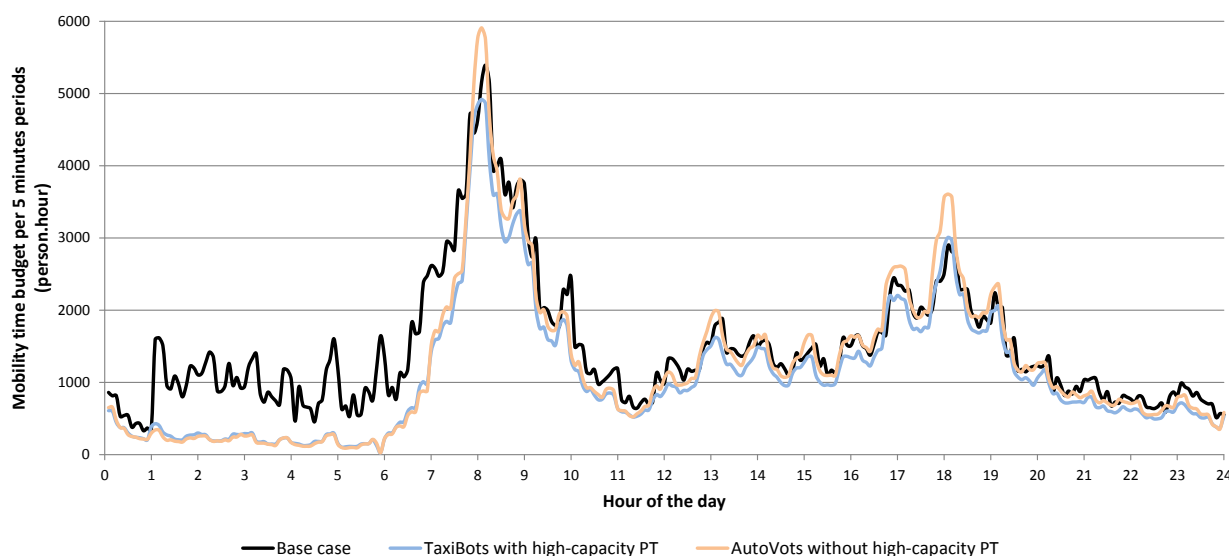(in periods of five minute over the course of a weekday)

Figure 8 shows the overall travel time budget (i.e. the amount of time spent travelling), logged for periods of five minutes over the course of a weekday. It shows that both ride-sharing systems with TaxiBots and car-sharing systems with AutoVots reduce the amount of travel time compared to the base case, with the exception of peak hour AutoVot trips. For TaxiBots complemented with high-capacity public transport, the total daily travel time required to deliver all trips is reduced by 30%. For AutoVots without high-capacity public transport, this reduction is lower (-18%). Although both TaxiBots and AutoVots save time on average, this is not necessarily true for all travellers and all times. Peak-hour AutoVot travel, for instance, will take more time than in the base case. For car drivers, however, switching to an AutoVot (without high-capacity public transport available) reduces the total door-to-door travel time by an average of 2%, assuming they spent three minutes per ride searching for parking space with their own cars. Car drivers who switch to ride sharing with TaxiBots with high-capacity public transport, on the other hand, would increase total door-to-door travel times on average by 8%.

Distribution of detour distance and time values

As noted earlier, ride sharing in the TaxiBot system implies acceptance of slight deviations from the shortest travel path for each passenger. The statistical distributions of the distance and time required for these detours are presented in Figure 9. The curve shows that the delays introduced by the detours are small, with less variability for the time detour. This suggests that additional travel time may not be a concern for ride-sharers in TaxiBots even during peak hours.

Figure 9. **Statistical distribution of detour for the TaxiBot scenario with high-capacity public transport**
(Top=distance, bottom=time; probability density function)
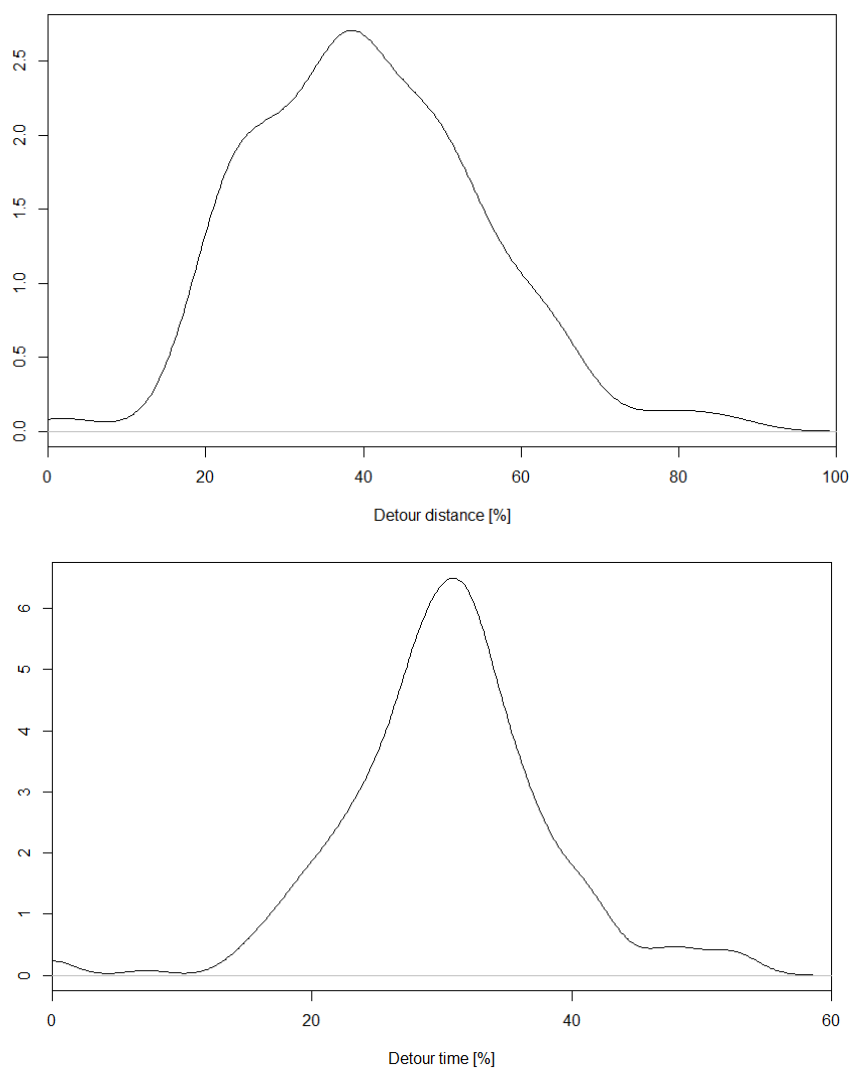
## Occupancy levels and efficiency of matching

We investigated the efficiency of matching users for shared rides, measured by average number of occupants, and how demand intensity and its patterns influence performance of TaxiBot ride sharing. Analysis of the average party size occupying TaxiBots throughout the day (Figure 10) reveals that peaks in demand in the morning and early evening enhance the ability to form larger ride-sharing pools. Even in periods of low average demand such as late evening or early morning, concentrated demand in popular

geographic areas - neighbourhoods with a high number of bars and clubs, for instance - still produce satisfactory average occupancy. Significantly, the concentration of boarding at origin, which prevents large waiting times (greater than five minutes in the tested scenarios), proved to be a more decisive component to ensure high occupancy than the density of alighting locations.

The overall daily elasticity of party size with respect to the aggregate city demand was 1.07. This means that by each 1% increase in overall demand the average party size is expected to increase 1.07%, indicating the elastic behaviour of this relation.

Figure 10. **Matching efficiency for shared rides in TaxiBot scenario with high-capacity public transport**
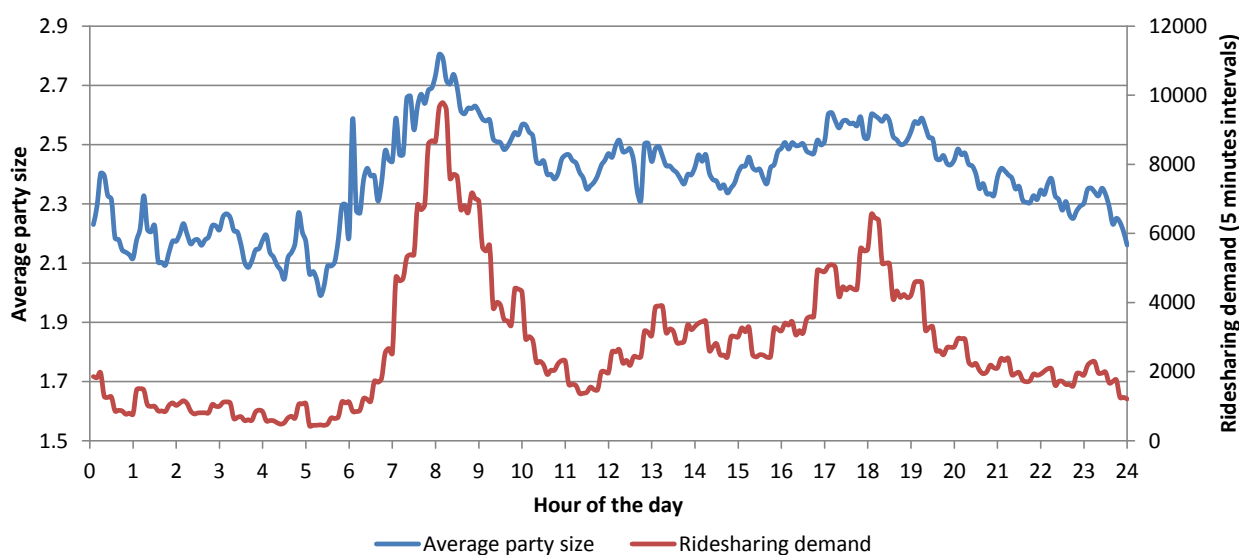
This study did not look at the costs of shared self-driving urban mobility, nor did it assess the costs to travellers of the base case trips. As such, the model is agnostic with respect to overall welfare impacts of the scenarios examined. In reality, travel demand is of course linked to costs, and further work could examine the extent to which savings in travel time and other potential reductions of costs for consumers might translate into additional travel demand. Part of that analysis could be the extent to which added demand might erode some of the potential benefits of the deployment of shared self-driving urban mobility systems. Our discussion also did not touch upon labour issues, which are likely to be significant, nor upon questions of equity that could be raised by such systems and that will have to be examined in further work.

5. Policy insights

Transport policies can influence the type and size of the car fleet, the mix between public transport and shared vehicles and, ultimately, the amount of car travel, congestion and emissions in a city. For small and medium-sized cities, it is conceivable that a shared fleet of self-driving vehicles could completely obviate the need for traditional public transport.

Shared vehicle fleets free up a significant amount of space both on and off-street. However, prior experience indicates that this space must be pro-actively managed in order to lock in benefits. Management strategies could include restricting access to this space by allocating it to commercial or recreational uses, delivery bays, bicycle tracks or enlarging sidewalks. Freed-up space in off-street parking could be used for logistics distribution centres.

Despite increases in overall levels of car travel, the deployment of large-scale self-driving vehicle fleets will likely reduce crashes and crash severity. Environmental impacts are tied to per-kilometre emissions and thus will be dependent on the degree to which shared self-driving car fleets employ more fuel-efficient and less polluting technologies.

The deployment of shared self-driving car fleets in an urban context will directly compete with the way in which taxi and public transport services are currently organised. These fleets might effectively become a new form of low capacity, high quality public transport. Labour issues will likely be significant, but there is no reason why public transport operators or taxi companies could not take an active role in delivering these services. This could alleviate these issues to a certain extent. Governance of transport services, including concession rules and arrangements, will need to adapt.

The drastic reduction in the number of cars resulting from a shift to shared self-driving cars will significantly impact car manufacturers' business models. New services will develop under these conditions, but it is unclear who will manage them and how they will be monetised. The role of authorities, both regulatory and fiscal, will be important in guiding developments or potentially maintaining market barriers.

Under all of our scenarios, vehicles are used much more intensely than before. From currently 50 minutes and 30 kilometres, daily use would increase to 12 hours and nearly 200 kilometres. This increase in use will require different car models than are currently on the market today. It will also mean shorter lifecycles and thus a quicker adoption of new, cleaner technologies. Shared use will also require different and much more robust interior fittings, although weight savings could potentially accompany a reduction of crash risk. Innovative maintenance programmes could be part of the monetisation package developed for these services.

Overall, vehicle travel will be higher in all fleet-mixing scenarios and vehicle numbers will increase in 3 out of 4 of our peak hour scenarios. It is likely that improved traffic flow could mitigate congestion up to a point. In the most extreme scenarios, however, it may be difficult to make a public policy case for self-driving fleets alone (without high-capacity public transport) based solely on space and congestion benefits. Nonetheless, even in mixing scenarios, these fleets could represent a cost-effective alternative to traditional forms of public transport if the impacts of additional travel are mitigated. "All in" deployment of shared self-driving fleets may be easier in circumscribed areas such as business parks, campuses, islands, as well as in cities with low motorisation rates.

Bibliography

L.D. Burns, W.C. Jordon, B.A. Scarborough (2013), "Transforming personal mobility", Earth Island Institute, Columbia University, 2013.

Câmara Municipal de Lisboa (2005) Lisboa: O desafio da mobilidade. Câmara Municipal de Lisboa - Licenciamento Urbanístico e Planeamento Urbano.

D. J. Fagnant, 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, March 2014.

L. M. Martinez, G. H. A. Correia, J. M. Viegas (2015), "An agent-based simulation model to assess the impacts of introducing a shared-taxi system: an application to Lisbon (Portugal)", *Journal of Advanced Transportation*, 49(3), April 2015.

P. Santi, G. Resta, M. Szell, S. Sobolevsky, S.H. Strogatz, and C. Ratti (2014), Quantifying the benefits of vehicle pooling with shareability networks, *Proceedings of the National Academy of Sciences of the United States of America (PNAS)*, 111(37), September, 2014.

K. Spieser, K. Ballantyne , K. Treleaven, R. Zhang, E. Frazzoli, D. Morton, M. Pavone (2014), *Toward a systemic approach to the design and evaluation of automated mobility-on-demand systems: A case study in Singapore. MIT Open Access Articles* (Forthcoming in Road Vehicle Automation, Springer Lecture Notes in Mobility series), 2014.

J.M. Viegas, L.M. Martínez (2010): "Generating the universe of urban trips from a mobility survey sample with minimum recourse to behavioural assumptions", in: *Proceedings of the 12th World Conference on Transport Research*, Lisbon.

J. Zachariah, J. Gao, A. Kornhauser, T. Mufti (2013): "Uncongested mobility for all: A proposal for an area-wide autonomous taxi system in New Jersey", in: *Proceedings of Transportation Research Board Annual Meeting*, Washington D.C.

Urban Mobility System Upgrade
How shared self-driving cars could change city traffic

What if all trips in a city were carried out by a fleet of self-driving cars shared by users? This study explores the potential outcomes of such a radical upgrade in an urban mobility system. It concludes that up to 9 out of 10 conventional cars could become redundant under certain circumstances. Vast amounts of public space would be freed for other uses in such a scenario. However, the total volume of travel increases in most scenarios and the net benefit of such an urban mobility system upgrade decisively depends on the choice of vehicle type, the level of penetration and the availability of high-capacity public transport to complement the shared self-driving car fleet.

The work for this report was carried out in the context of a project initiated and funded by the International Transport Forum's Corporate Partnership Board (CPB). CPB projects are designed to enrich policy discussion with a business perspective. Led by the ITF, work is carried out in a collaborative fashion in working groups consisting of CPB member companies, external experts and ITF researchers.