SHIFT

THE AUTONOMOUS VEHICLE DEPLOYMENT REPORT

Contents

Foreword

by Dr Graeme Smith, Senior VP of External Affairs, Oxbotica

As a society, we are beginning a technological shift in the way we travel and move goods, which is so significant that it could almost be compared to the invention and roll-out of the computer. Autonomy, whether we're talking purely in terms of transportation or any of its various other use cases, is set to transform the world we live in.

It has the potential to enhance our safety, reduce carbon emissions and enrich our quality of life. If any vehicle, at any time, in any place can become autonomous, there really are no limits, the opportunities are endless, and we all stand to benefit.

In order to achieve this, it is vital to define a path from proof-of-concept connected autonomous vehicle (CAV) demos to larger-scale deployments – and not just in the vehicles themselves. Humans have been making vehicles for over a hundred years, and the fundamentals of vehicles are firmly established. We know how to make them cheaply and quickly, and they are evolving to be ever more sophisticated, capable and safe.

The knowledge gap is how CAVs may impact the area where they are deployed. Questions such as "what is required from the external environment, for example road infrastructure or charging points?" and "how will the travel behaviour of citizens be expected to change after CAVs are deployed?" remain unanswered.

And that's why SHIFT is so important and unique. It has helped push the boundaries of autonomy, not just in the technology that drives the vehicles but in developing and defining the nuts and bolts of real-world deployment. Successful autonomous deployment requires a global autonomy software platform that uses smart design and modern coding techniques to create future-proofed and open APIs between components.

The SHIFT consortium, supported partly from a £1.58 million grant awarded by the Centre for Connected and Autonomous Vehicles (CCAV) and delivered through Innovate UK, provides the basis for supporting proof-of-concept scale up from small CAV demos to large scale commercial deployments. By doing this, we can collaborate with an ecosystem of partners and contributors big enough to address the massive opportunities in the AV sector.

The outputs of SHIFT put us on the road to a world where any third party – from OEMs through to owners and operators – would be able to deploy CAV services in harmony with an urban environment in a way that suits their business model. This is the key to unlocking true Universal Autonomy.

Executive Summary

SHIFT is a collaborative project between Oxbotica, Imperial College London, and Transport for London (TfL) to start to model the impact of autonomous vehicles (AVs) on various urban environments and develop guidance for deploying fleets safely and efficiently.

Over the last two years, Imperial College London and TfL have used a combination of transport analysis tools to start to develop the capabilities to model the impact of AVs on city transport, congestion and the mode-shift away from vehicle ownership. Alongside this, Oxbotica has focused on the building blocks of wider scale autonomous deployment, including the vehicle build and operator training, helping to define a standard that can be efficiently replicated.

The outcome is the 'SHIFT Autonomous Vehicle Deployment Report', which will inform autonomous fleet operators as they begin to deploy autonomy services, and provides guidance on the impact AVs could have on existing infrastructure as well as how to regulate fleets. The work completed is a key step in taking proof-of-concept autonomous demonstrations to larger-scale service deployments in the UK.

The project has identified the key challenges fleet operators rolling out AVs might face in urban areas, such as coping with vehicle downtime, charging infrastructure, interaction with existing public transport, the optimal fleet size and reducing vehicle mileage by anticipating future demand.

Three urban use cases have been modelled in the SHIFT Autonomous Vehicle Deployment Report, including a central city example, an inner-city deployment and an outer city scenario. For each case study, the variations in ride-sharing probability and fleet size have been modelled to enable fleet operators to make accurate decisions during deployment planning depending on their requirements. Deploying the correct fleet size that matches demand will allow operators to reduce customer waiting time, optimise routing of vehicles to create less traffic and reduce emissions.

The SHIFT Autonomous Vehicle Deployment Report has also created the basis for fleet health monitoring to improve safety. The creation of a dashboard of metrics is necessary to allow fleet operators to keep vehicles in top condition by tracking trends to predict any issues before they develop into a safety issue. Similarly, a data infrastructure has been produced as part of SHIFT to standardise Oxbotica's collecting, recording and downloading of vehicle logs.

The project also provides the basis for building and commissioning AVs efficiently and safely. A comprehensive and repeatable build order has been produced during SHIFT, detailing the hardware and engineering required to make a vehicle AV-ready in a way that maximises trust in the technology.

Safety is further improved through the project with the creation of training standards for drivers and safety drivers in the vehicle when operations are handed over to a third party for deployment. Test drivers are a key part of the deployment, therefore ensuring consistent training processes are in place is vital to both the safety and efficiency of AV deployment roll-out.

The completion of the SHIFT project and the subsequent publication of the SHIFT Autonomous Vehicle Deployment Report answers questions on autonomous fleet deployment in a way no other consortium or research project has before.

From training and infrastructure considerations through to simulation and data modelling, SHIFT lays down the key considerations and impacts of deploying autonomous vehicles at scale in a city or urban environment.

What is SHIFT?

SHIFT is a collaborative project between Oxbotica, Imperial College London and Transport for London (TfL) supported partly from a £1.58 million grant awarded by the Centre for Connected and Autonomous Vehicles (CCAV) and delivered through Innovate UK. It aimed to develop models and guidelines that will facilitate the appropriate deployment of autonomous vehicles (AVs) safely and efficiently in urban environments. The project partners used transport analysis tools to start to model the impact of AVs on city transport, congestion and the mode-shift away from vehicle ownership.

We describe SHIFT as defining the key considerations and impacts of deploying AVs in urban environments.

The SHIFT Partners

oxbotica

The project has been led by Oxbotica, the global software company deploying autonomy around the world. Oxbotica provides an autonomy software platform designed to bring autonomy solutions to clients and partners in the shape they need it, and using the sensors and physical vehicles that suit them.

Oxbotica's software is being used today all around the world - on streets, in quarries, around runways, and down mines - helping to drive efficiencies, increase safety and improve production. During SHIFT, Oxbotica has been responsible for overseeing the vehicle build and the development of the fleet management system, which can be scaled as required.

Imperial College London

TSL Transport Systems & Logistics Laboratory

The Transport Systems and Logistics Laboratory (TSL) at Imperial College London focuses on the study of networks, optimisation methods and multi-agent systems, as well as their applications in autonomous transport systems, urban infrastructure and logistics. As part of the SHIFT project, TSL developed a digital twin of the autonomous ride-sharing fleet deployments, underpinned by a bespoke urban mobility model that was developed in-house.

Through an extensive set of simulation campaigns that took place over the SHIFT project's duration, TSL researchers studied the decision processes that could affect the demand for an autonomous ride-sharing fleet, including optimal fleet assignment algorithms, energy management simulations and road user predictions.

For more information on the analysis carried out by TSL click here.

TRANSPORT FOR LONDON
EVERY JOURNEY MATTERS

Modellers and analysts from TfL were involved in the modelling element of the work that the consortium completed. Using an understanding of London's road network and the demand for travel in the city, TfL was able to develop an approach to explore the potential implications of connected autonomous vehicles (CAVs) on different areas within London.

The modelling can support the central goal of the Mayor of London's Transport Strategy (MTS) to ensure that by 2041, 80 percent of all trips in London are to be made on foot, by cycle or using public transport and that London's air quality is improved. The Strategy is underpinned by the Healthy Streets Approach, which provides the framework for putting human health and experience at the heart of planning the city. A further vital element is the Mayor's ambition for Vision Zero – that by 2041, all deaths and serious injuries will be eliminated from London's transport network. So, it is important for TfL to be able to model the potential implications of CAVs on London's roads and the impacts there might be on these objectives. More information on TfL's thinking on CAVs and its guidance for London trials can be found here.

As part of the project, TfL has shared current trip and mode share data with the Imperial College London modelling team to help understand how the deployment of CAVs may vary across different areas of London. The model that was developed can explore the potential impact of different CAV scenarios on the road network and on mode shift.

The Autonomous Vehicle Deployment Report

The Delos Urban Mobility Model

Developed by the Transport Systems and Logistics Laboratory at Imperial College London

The simulation platform used in the SHIFT project consists of three key components:

- a customer behaviour model, capturing user preferences and choice logic
- a fleet operations controller, focusing on vehicle logic and platform management
- a digital twin of the transport network where customers and vehicles interact

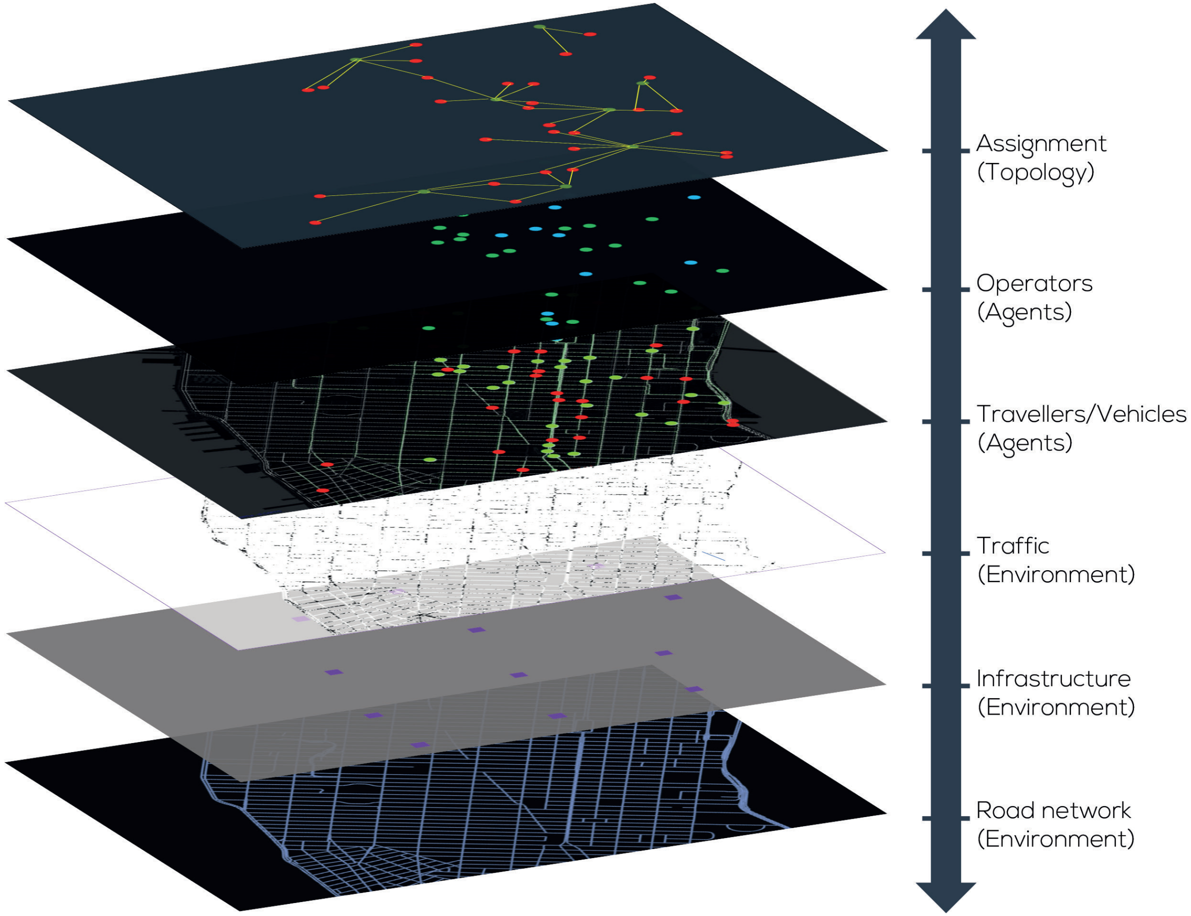

Agents

The Delos Urban Mobility platform uses three types of agents to represent a digital twin of the autonomous vehicle (AV) fleet deployment: travellers, vehicles, and operators. Using these, we can represent hundreds of thousands of interacting entities, alongside the logic rules that dictate every aspect of their behaviour.

Travellers in our model have the ability to plan their trips using a wide range of transportation modes, with shared AVs being one of the options. In deciding which mode of transport they will choose, travellers consider aspects such as affordability, waiting times, travel times and convenience. The model is dynamic, and as such, should any congestion or operational delays occur, they are free to choose an alternative mode of transport.

Vehicle agents can serve individuals or groups of customers by providing door-to-door transport between their origin and destination points. The level of service offered to customers, and revenue from collected fares is influenced by platform policies, the overall service configuration, and the choice of assignment or pricing strategies implemented by fleet operators.

Operating Regime

Each agent in the model is aware of their neighbours' actions. A typical traveller in the model would, in the first instance, submit a trip request to the fleet operator. The operator would in turn seek to identify a vehicle that could serve the customer.

Once an available vehicle has been found, the platform would inform the traveller how long it will take for the vehicle to arrive at their location, what is the expected travel time of the journey, and how much it is going to cost. If the traveller accepts the above (as opposed to opting for an alternative mode of travel), the vehicle will be instructed to proceed with the journey.

Network Topology

The network topology represents the spatial characteristics of the simulated environment. The Delos Urban Mobility platform is capable of generating a full digital twin of the deployment region, using data extracted from openly available data (e.g. OpenStreetMap, public transport timetables), which can be further enhanced using statistical information that describes travel demand and other aspects of customer behaviour. An additional layer in the model is used to represent the infrastructure components that will support fleet operations, such as on-street and off-street parking provision (used by the vehicles while they are idle), refuelling stations and EV charging depots across the network.

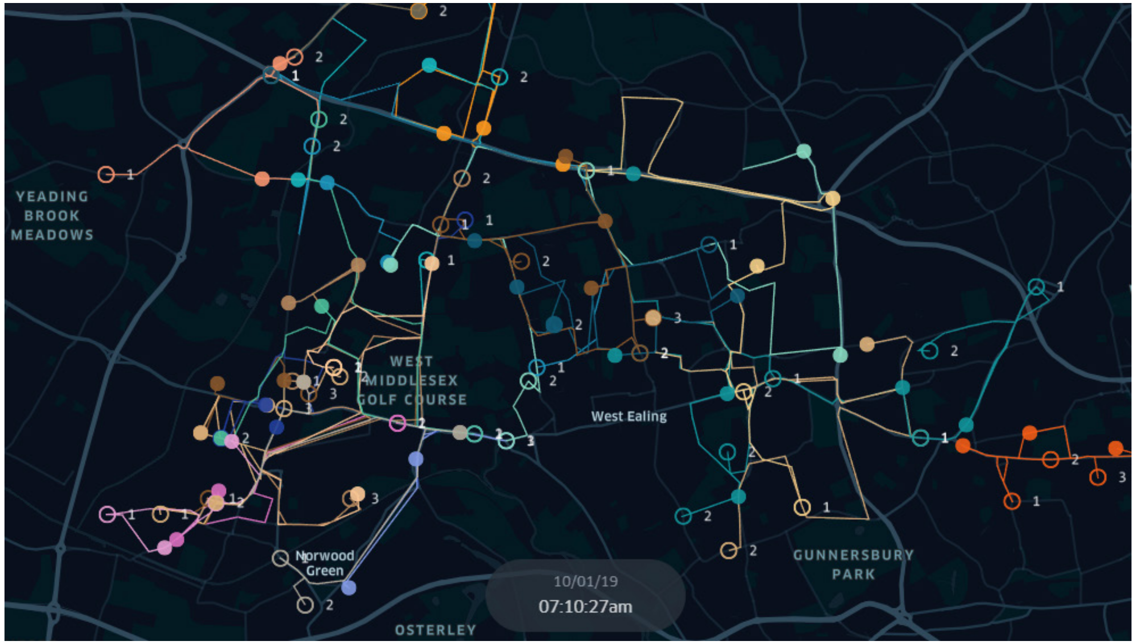

Case Studies

Three London regions were selected as test cases to evaluate the deployment of autonomous fleet services in different environments: a central, inner and outer environment. For the purposes of this study, we tested the fleet control algorithms by varying the following operational parameters:

- **Ride-Sharing Acceptance (0% - 100%)** – indicates the likelihood that a customer is willing to share the vehicle with other customers that are travelling in the same direction. Analysis suggests that this parameter will increase trip durations and detour times for customers, as well as increase the average energy expenditure per vehicle, but will allow smaller fleet sizes to cope with similar demand compared to non-sharing approaches.
- **Fleet Size (50 - 250)** – determines the number of vehicles used in a fleet, which should reduce the customer abort rate as its value increases. However, if the fleet size exceeds a certain threshold, the number of idle vehicles in the network is expected to grow, as will the operational cost of running the service.
- **Vehicle Capacity (4 - 6)** – determines how many passengers can be travelling with the vehicle at any time (either as a single occupant or as multiple groups of customers). In the cases where shared rides are allowed, an increase of this parameter would have a similar effect as a potential increase in fleet sizes.
- **Vehicle propulsion** – the fleet is assumed to be homogenous, so all vehicles will share the same engine characteristics (either diesel or electric). Diesel-powered vehicles will emit NOx and Particulate Matter (which can be tracked by our model). On the other hand, electric vehicles have a smaller range and need to be taken offline for recharging.

Performance Analysis and Indicators

A range of performance indicators was used to obtain an in-depth understanding of how the platform performs under a given operating regime. The metrics used can consider different aspects of customer behaviour, fleet performance and the utilisation of supporting infrastructure.

An interactive dashboard that contains a detailed breakdown of the results obtained by this study is available here.

1. Trip Sharing

Number of distinct customer groups sharing the same ride

Customers have the option of sharing a ride with other travellers, in exchange for a reduced trip fee. On such occasions, the operator would attempt to match multiple customers that are travelling in the same direction.

Our analysis indicates that there is no clear relationship between customer sharing and the size of the fleet. However, a clear correlation exists between vehicle capacity and the incidence of shared rides.

Shared rides per customer

This indicator monitors the number of shared trips that a single customer will participate in, as the vehicles make their way through their region and pick up or drop off different groups of customers. Analysis indicates that smaller fleet sizes will result in a larger proportion of shared rides, as the platform would have a reduced stock of vehicles that could absorb travel demand.

The two sharing indicators can be used to determine the ideal composition of the fleet to be deployed in a region. A higher incidence of shared trips may be indicative of a need to increase the fleet size.

Customers that participate in shared trips may be subject to longer travel durations, as a result of detours that would need to take place as vehicles are tasked with collecting additional customers. The Delos Urban Mobility model tracks detour times in order to quantify the "inconvenience" of a shared ride to individual customers.

Proportion of shared trips

Irrespective of customer appetite for trip sharing, it might be difficult for the platform to match customers due to other aspects of the fleet management regime or due to the spatial characteristics of the deployment region. The difference between customer appetite for sharing and the proportion of shared trips that were realised can be used to quantify the effectiveness of the trip matching process.

2. Fleet utilisation

Vehicle idle time

Vehicles are recorded as idle when they are stationary (parked), or relocating to a different region, in anticipation of trip requests that have not yet materialised. Higher idle times would indicate that the fleet is not utilised efficiently, and there potentially exists scope to reduce the number of vehicles that are deployed.

Vehicle distance travelled

This metric tracks the total distance travelled by vehicles across all stages of operation (customer trips, pick-up journeys, relocation, refuelling). It can be used in conjunction with the idle time metric to determine the optimal fleet size for a particular deployment. A decrease in the average distance travelled as more vehicles are added to the fleet, might also be indicative of poorer fleet utilisation.

Vehicle speed

Average speeds are dependent upon the choice of route, and in turn influence trip times and energy expenditure. It can be improved using better vehicle routing algorithms and through careful consideration of traffic congestion data.

Vehicles idle

This metric focuses on the number of vehicles that might be idle at any point in time over the course of daily operations. A large number of idle vehicles would suggest that the size of the fleet can be reduced without negatively affecting the level of service that is provided.

A small proportion of idle vehicles in the fleet is desirable, as it would ensure that there exists availability to promptly serve customers once a request is lodged with the platform.

Recharging activity

All vehicle types have a limited travel range, which depends on the mode of propulsion, the size of their battery or fuel tank, and the choice of fleet energy management regime. A range of fleet management algorithms were developed by TSL that would seek to schedule refuelling or recharging sessions at periods of reduced travel demand.

3. Customer operations

Trip duration

Measures the time that customers spend aboard a vehicle, while being transported between their pickup location and their destination. Even though shorter trip durations may suggest that the fleet is managed effectively, it could be indicative of an inability to accommodate longer (and potentially more profitable) journeys.

Detour time

Customer detours are observed in the majority of shared rides, as vehicles need to veer slightly off-course in order to collect another customer. From a ride-hailing viewpoint, this value should be minimised in order to reduce any delays that might be experienced by the customers.

Pick-up time

This measure focuses on the waiting time between the assignment of a vehicle to a customer and the actual time that the pick-up occurs. If this waiting time is expected to exceed a threshold specified by the customer, the trip request is withdrawn, as the customer explores alternative options. Lower values of this indicator are indicative of good fleet management practices. Pick-up times can be reduced through the use of more sophisticated matching algorithms and the deployment of more vehicles in the region.

Queuing time

This indicator focuses on the amount of time required by the platform to identify a vehicle that could accommodate a customer request. Queuing times are longer in periods of peak demand or in deployments with insufficient vehicle numbers.

Aborted customers

Customers can still withdraw their trip request after a vehicle assignment has been accepted. Aborted trips would occur while they are still waiting for their pick-up and are indicative of unforeseen congestion levels, poor energy management or vehicle failure. Once a trip has been aborted, the customer would continue their journey with an alternative mode of transport.

User throughput

The average number of users that are travelling at any point in time depends both on the current levels of travel demand, as well as on the ability of the fleet to absorb it. A low user throughput during periods of high demand would be indicative of poor fleet management practices.

4. Energy usage

Mode comparison

As part of this study we carried out a comparison of energy consumption patterns across the various modes of transport that are included in the Delos Urban Mobility model. The energy comparisons shown in the charts below are as follows:

An energy consumption analysis was carried out across three deployment scenarios, involving normal, high and very high travel demand patterns.

The high enery consumption in trips using CAVs reflects the increased complexity of fleet management operations.

While these are preliminary results, they represent a baseline level that can be reduced significantly using algorithims.

Rail: Travellers using fixed track modes of transport (Tube, Train, DLR) instead of CAVs.

Bus: Travellers using buses instead of electric CAVs.

Private: Travellers using private vehicles instead of CAVs.

CAV: Travellers using CAVs.

AVERAGE ENERGY: NORMAL DEMAND

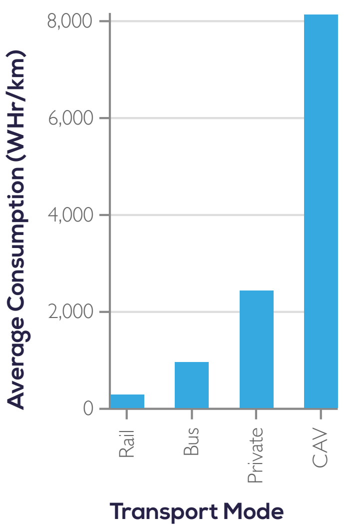

AVERAGE ENERGY: HIGH DEMAND

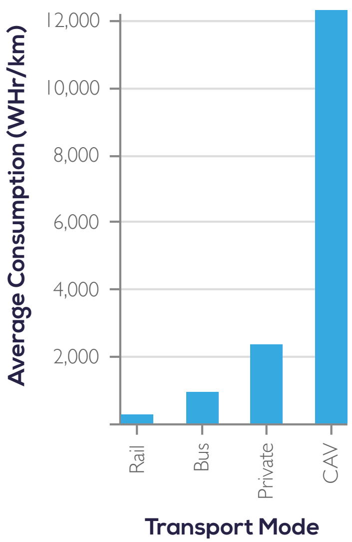

AVERAGE ENERGY: VERY HIGH DEMAND

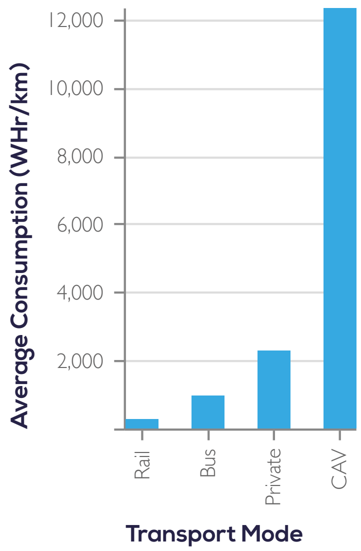

4. Energy usage

The following indicators were used to quantify the performance of the fleet with respect to energy consumption:

Trip energy consumption

Focusing on individual trips, and in accordance with the mode of vehicle propulsion and the speed of travel.

Vehicle energy consumption

Incorporating relocation, pick-up and service travel. This metric is indicative of the effectiveness of the fleet management algorithms that are used by the platform.

5. Platform Costs

Vehicle operating costs

This indicator considers the upfront costs for fleet procurement, insurance, maintenance and operation costs (as a function of vehicle activity).

Per trip costs

Trip costs vary in accordance to the size of the fleet, distance travelled, and the times of operation. Higher levels of trip sharing among customers would correspond to increased trip costs, as a result of longer trip times due to detours. However, it must be noted that trip sharing would potentially result in a higher user throughput, and overall revenues.

6. MoTiON Model development by TfL

TfL's own multi-modal demand model; Model of Travel in LondON (MoTiON) is now ready to test various connected autonomous vehicle (CAV) business models (e.g. personally owned vehicles, ride hailing services for individuals and shared ride hailing services carrying multiple independent journeys) by varying fare structures, sharing behaviour, and the amount of diversion and empty running.

Working with outputs from the Delos Urban Mobility model gives the opportunity to define some of the parameters in MoTiON's CAV functionality not previously possible, for example key performance indicators such as individual waiting times, journey times and detour time (for shared trips).

7. Optimal Fleet Operation

The multitude of performance indicators that are tracked by the Delos Urban Mobility model allows users to study and fine-tune a wide range of parameters that can influence the operational performance. This example focuses on the selection of an optimal fleet size, which is a key stage in the deployment of a ridesharing platform.

As an example, while a reduction in pick-up times would increase customer satisfaction, it should not be accompanied by a reduction on the overall throughput, as this would suggest that the platform rejects trip requests for the benefit of few customers.

To account for the complex dynamics of ridesharing operations, the Delos Urban Mobility platform uses a multi-objective approach in order to select the optimal size of the fleet, as it seeks to strike a balance between several conflicting priorities.

Aborted trip requests

Minimisation of aborted trip requests (to improve customer satisfaction).

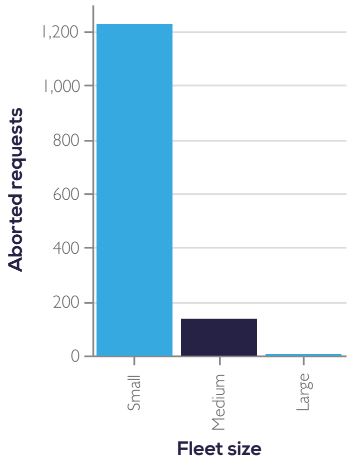

Energy consumption

Minimisation of trip energy consumption (in accordance to sustainability objectives).

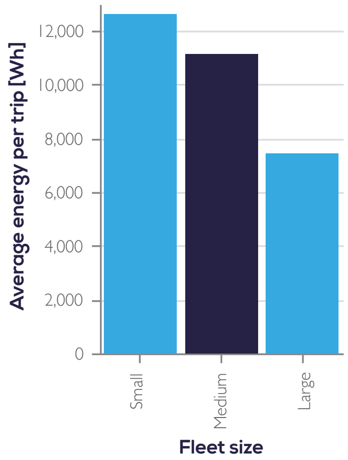

Customers served

Maximisation of the number of customers served (to increase revenue).

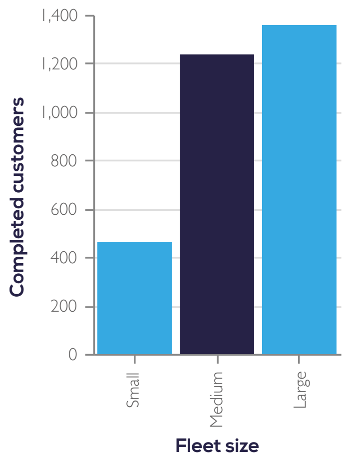

Vehicle operating costs

Minimisation of vehicle operating costs.

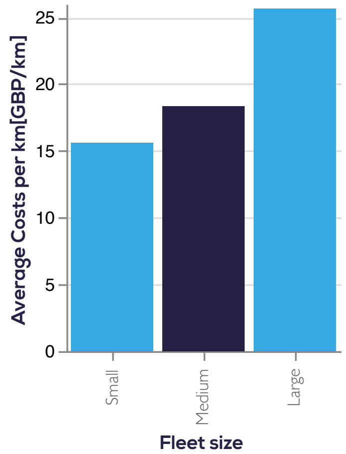

Customer detours

Minimisation of customer detours (to improve customer satisfaction).

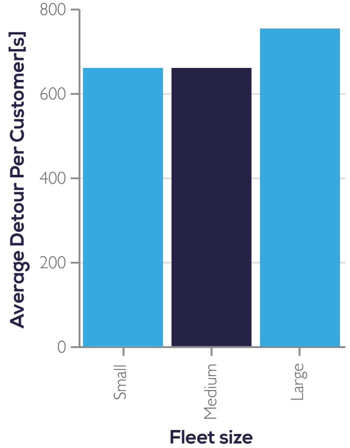

Idle vehicles

Minimisation of idle vehicles (to improve fleet utilisation).

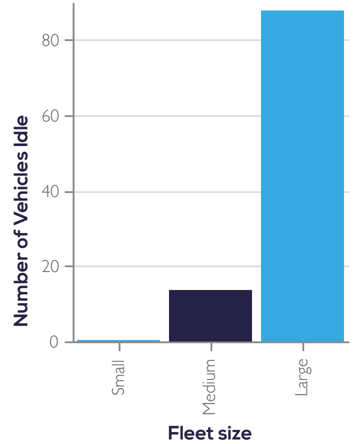

Trip durations

Minimisation of trip durations (to improve throughput and customer satisfaction).

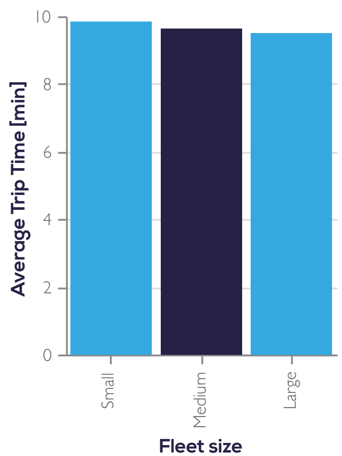

Customer waiting time

Minimisation of customer waiting time (to improve customer uptake and satisfaction).

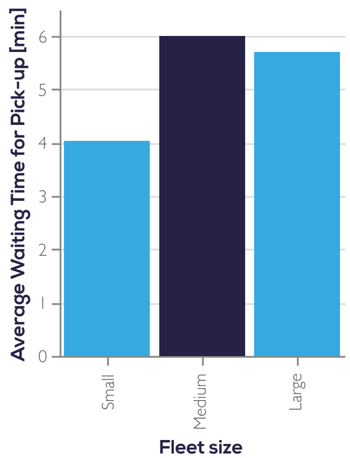

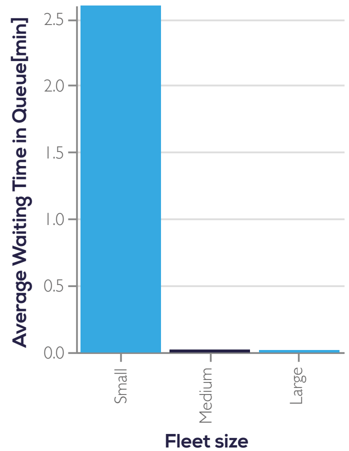

Small Fleet

A small fleet (50 vehicles) would not be able to absorb a large proportion of the customer base. Lack of vehicle availability would result in increased waiting times, which would in turn discourage customers from lodging trip requests with the platform in the long term. A negligible proportion of idle vehicles indicates that the service operates beyond capacity.

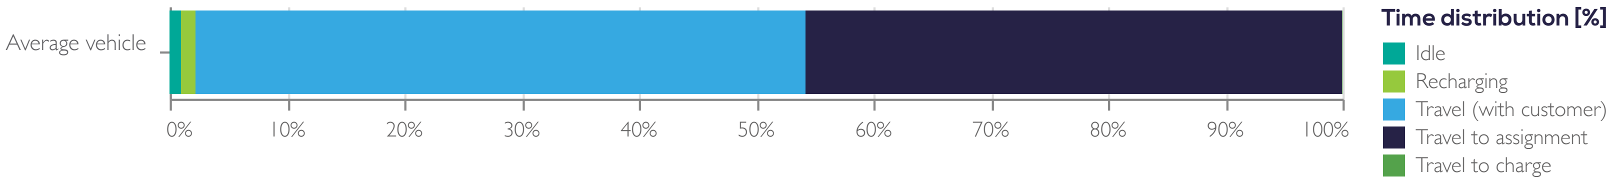

Medium Fleet

A medium-sized fleet (150 vehicles) is best suited to accommodate the demand scenarios in this particular deployment: most customers are served while detouring is kept at a low average value of 10 minutes. Higher operating costs are compensated by increased revenues, as the fleet is able to serve 1,000 more travellers compared to the previous scenario. A low number of idle vehicles suggests that there is still scope to increase the size of the fleet, or to invest in better fleet management algorithms.

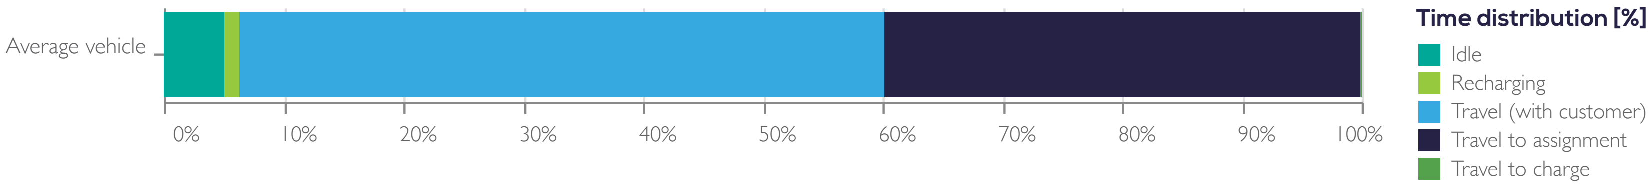

Large Fleet

A large fleet (220 vehicles) is seen to provide the best levels of customer service, reducing average trip times and wait times as well as servicing virtually all customer requests. However, this scheme results in poor fleet utilisation, as a large portion of the fleet (80+ vehicles) remains idle. Vehicle operating costs are increased by 40% compared to the previous scenario, resulting in reduced profits.

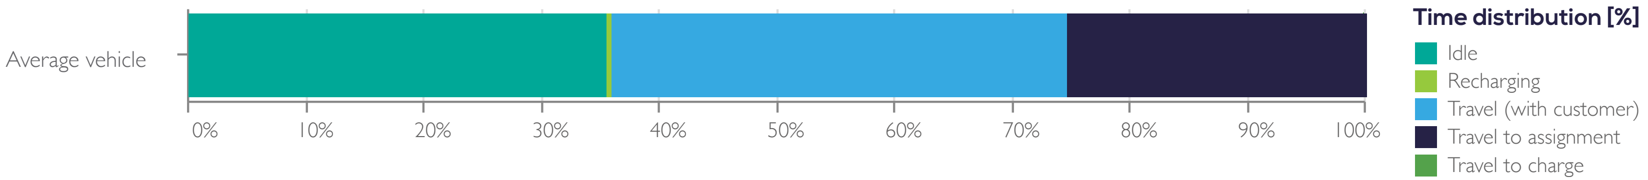

Fleet Management and Integration

Developed by Oxbotica and Imperial College London

The deployment of autonomous vehicles (AVs) requires a robust fleet management system to ensure the right vehicles are dispatched to the right user at the right time. Whenever a fleet is deployed in a new location, it is important to ensure the fleet operator can access the data, which requires onsite hardware to allow data ingress from vehicles to the infrastructure. The systems need to have high connectivity to support large data transfers while maintaining secure and robust channels. Once a vehicle is deployed, the operator will be required to monitor the vehicle's performance, using key performance indicators (KPIs) to track the data.

Efficient data management is a key element in successful usage of autonomous fleets. The sensor information collected from a single vehicle during mapping is in the range of terabytes. The challenge is of particular importance when the fleet is operated by a third party, and data needs to be stored, processed and transported without disruptions across multiple sites. Similarly, as fleets scale in size as the greater deployments are realised, the system must be capable of integrating multiple vehicles, from short-trip autonomous shuttles to longer-journeyed ride-sharing vehicles.

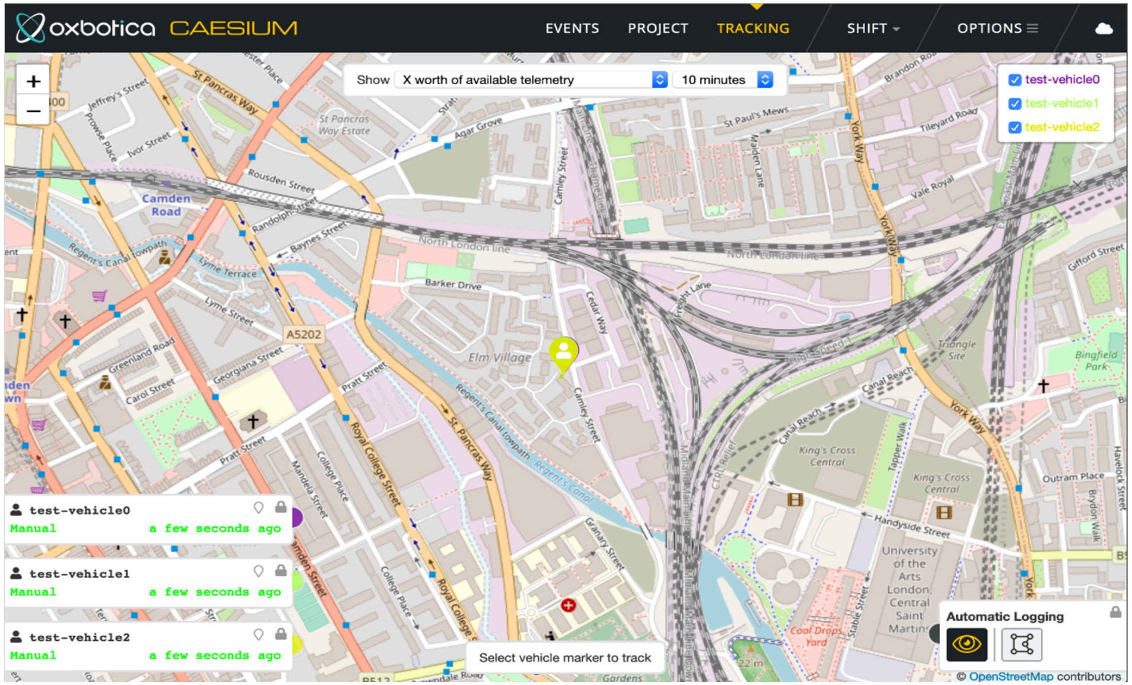

The development of an application programming interface (API), developed by Oxbotica and Imperial College London, is a key part of the SHIFT Autonomous Vehicle Deployment Report, that will allow a third-party fleet management system to be efficiently integrated with on-vehicle autonomous software. The development of this API means the fleet management system can dispatch vehicles and draw down data directly from vehicles to perform fleet and vehicle health checks, directly improving operational efficiency.

Fleet health and performance monitoring

SHIFT is the first project that has tackled large scale data sets from autonomous vehicles (AV) to create a dashboard of metrics for fleet management. The development of this dashboard by Oxbotica and Imperial College London is a key enabler for analysing the data and identifying which metrics are most important for fleet monitoring. This metric package is now at a stage that it could be shared with fleet operators of any scale, allowing them to track vital data across the deployment.

This can be used to enable real-time fleet tracking using key metrics. Tracking includes key performances over time, such as autonomous kilometres driven, vehicle performance to aid software development, key events such as network and CPU usage, and validation data such as steering calibration. The combination of this data set is vital for detecting any issues with individual vehicles as quickly as possible and allowing operators to make continual improvements to efficiency.

For example, one of the key metrics recorded and logged is lateral acceleration. Using accurate planned localisation readings and comparing them to the actual data downloaded from a vehicle, fleet operators can efficiently determine any problems with wheel alignment, tracking or tyre wear. Trends in the data will then allow fleet managers to predict and identify issues before they happen, to keep the fleet at 100% health and minimise unexpected downtime.

Similarly, when algorithm changes are made to AV software, the performance monitoring allows the real-world impact, for example smoothness of ride, to be addressed statistically rather than based on subjective feedback from a test driver. The data also allows comparison between the real driver and autonomous driving to help refine the ride experience.

Data infrastructure

In order to facilitate the fleet monitoring and performance metric gathering, SHIFT has had to develop a data infrastructure to collect, record and download the vehicle logs. The SHIFT solution is a repeatable data set-up that can be recreated by fleet operators in their own deployment.

The infrastructure, defined through a SHIFT API, allows logs to be extracted from the vehicle onto a storage device, transferred to the web app and then outputted into the monitoring dashboard.

Development of the data infrastructure requires on-road vehicles to log journeys in order to make data available for download. On-road activities were disrupted during the COVID-19 lockdown, yet development was able to continue with the creation of a test rig. Computer hardware similar to that used in-vehicle was utilised, providing a repeatable test environment. These simulations, created during the SHIFT project, have now become a vital part of the autonomous software development programme for Oxbotica.

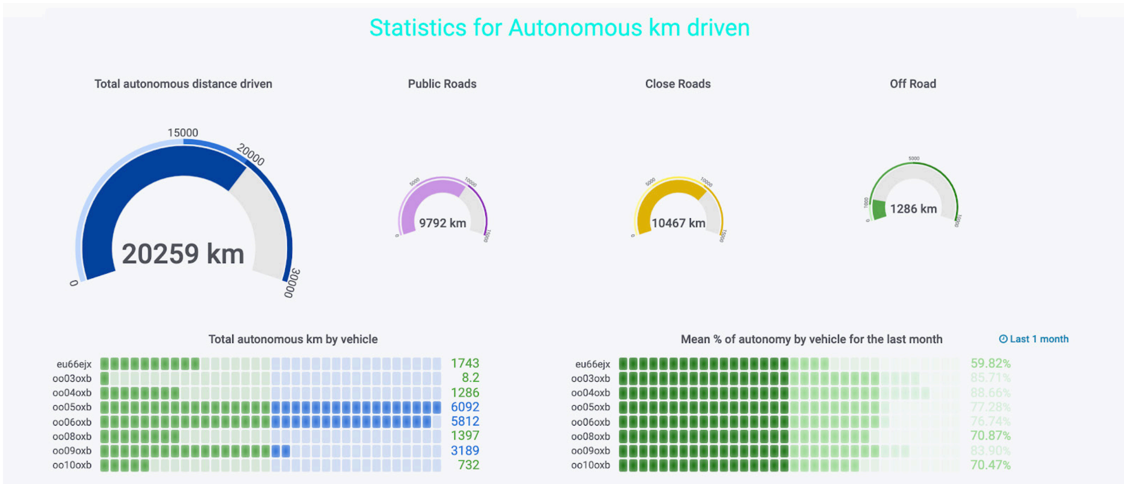

Safety

Developed by Oxbotica

Autonomous vehicles (AVs) have the potential to improve road safety by reducing collisions as well as improving mobility for groups who are often excluded by their inability to both access or drive conventional vehicles. From on-road scenarios to vehicle health and software redundancies, safety is an integral part of any AV deployment.

Operator training

The SHIFT project has helped Oxbotica define a consistent training process for those in the vehicle during public road demonstrations – both the Safety Driver behind the wheel and the Autonomous Control Systems Operator (ACSO) in the passenger seat. The role and required training of both the Safety Driver and ACSO are key parts of Oxbotica's safety case for public environment field trials of AVs.

For the first time, a full Safety Driver and ACSO training portfolio has been developed that can be used when operations are handed over to a third party as a part of commercial deployment. It enables inexperienced operators to gain the skills required before being measured and evaluated against strict criteria, to ensure safe operation and deployment.

As with many guides, the level of experience and training required varies depending on the environment and scenarios the operators will be deployed into. More complex activities require a higher level of expertise and experience.

As autonomous systems become more complex, defining the skills and requirements of the Safety Driver will become increasingly important for safe deployment. Alongside data metrics offloaded from the vehicle during testing, issue reporting and driver feedback from those in the vehicle is a key part of the development loop to ensure the vehicles deliver the safest and smoothest ride possible.

Vehicle Build

Developed by Oxbotica

In order to scale-up and move from a proof-of-concept R&D vehicle to a deployable fleet, a build specification is required that is efficient in both time and cost. The definition of a build schedule maximises trust in commercial deployments, from both the operator's and the customer's view. The work conducted during SHIFT will enable fleets to deploy safely, more efficiently, and at greater scale than ever before, with the ability to replicate the build process across multiple vehicles.

SHIFT has taken the outputs of DRIVEN (an earlier Oxbotica consortium project which deployed a fleet of six autonomous vehicles (AVs) onto the streets of East London) – such as what components to use, where to locate them on the vehicle for maximum efficiency, and acceptable sensor positional tolerances – and structured them into a build order. The developed instructions still focus on the same vehicle platform (Ford Mondeo Hybrid), but also define a template for the build specification format, with a detailed Bill-of-Materials (BOM), and instructions and assembly drawings for the build of components and assemblies.

Sensor type	Description	Purpose	Quantity
Camera (stereo)	Forward facing stereo camera	Localisation	1
Camera (stereo)	Rear facing stereo camera	Localisation	1
IMU	An inertial measurement unit (IMU) is an electronic device that measures and reports a body's specific force, angular rate, and sometimes the orientation of the body, using a combination of accelerometers, gyroscopes, and sometimes magnetometers	Localisation	1
Camera (monocular)	Forward facing cameras angled 45 degrees from direction of travel and in line with direction of travel. Mounted on sides of the roof bar and inside the vehicle	Perception	2
LiDAR	Main LiDAR mounted on top of the vehicle	Perception	1
LiDAR	Side LiDAR mounted at ends of the roof bar at an angle	Perception	2
LiDAR	Mounted to front grille	Perception	1
Radar	Mounted to front grille	Perception	1

Vehicle Build

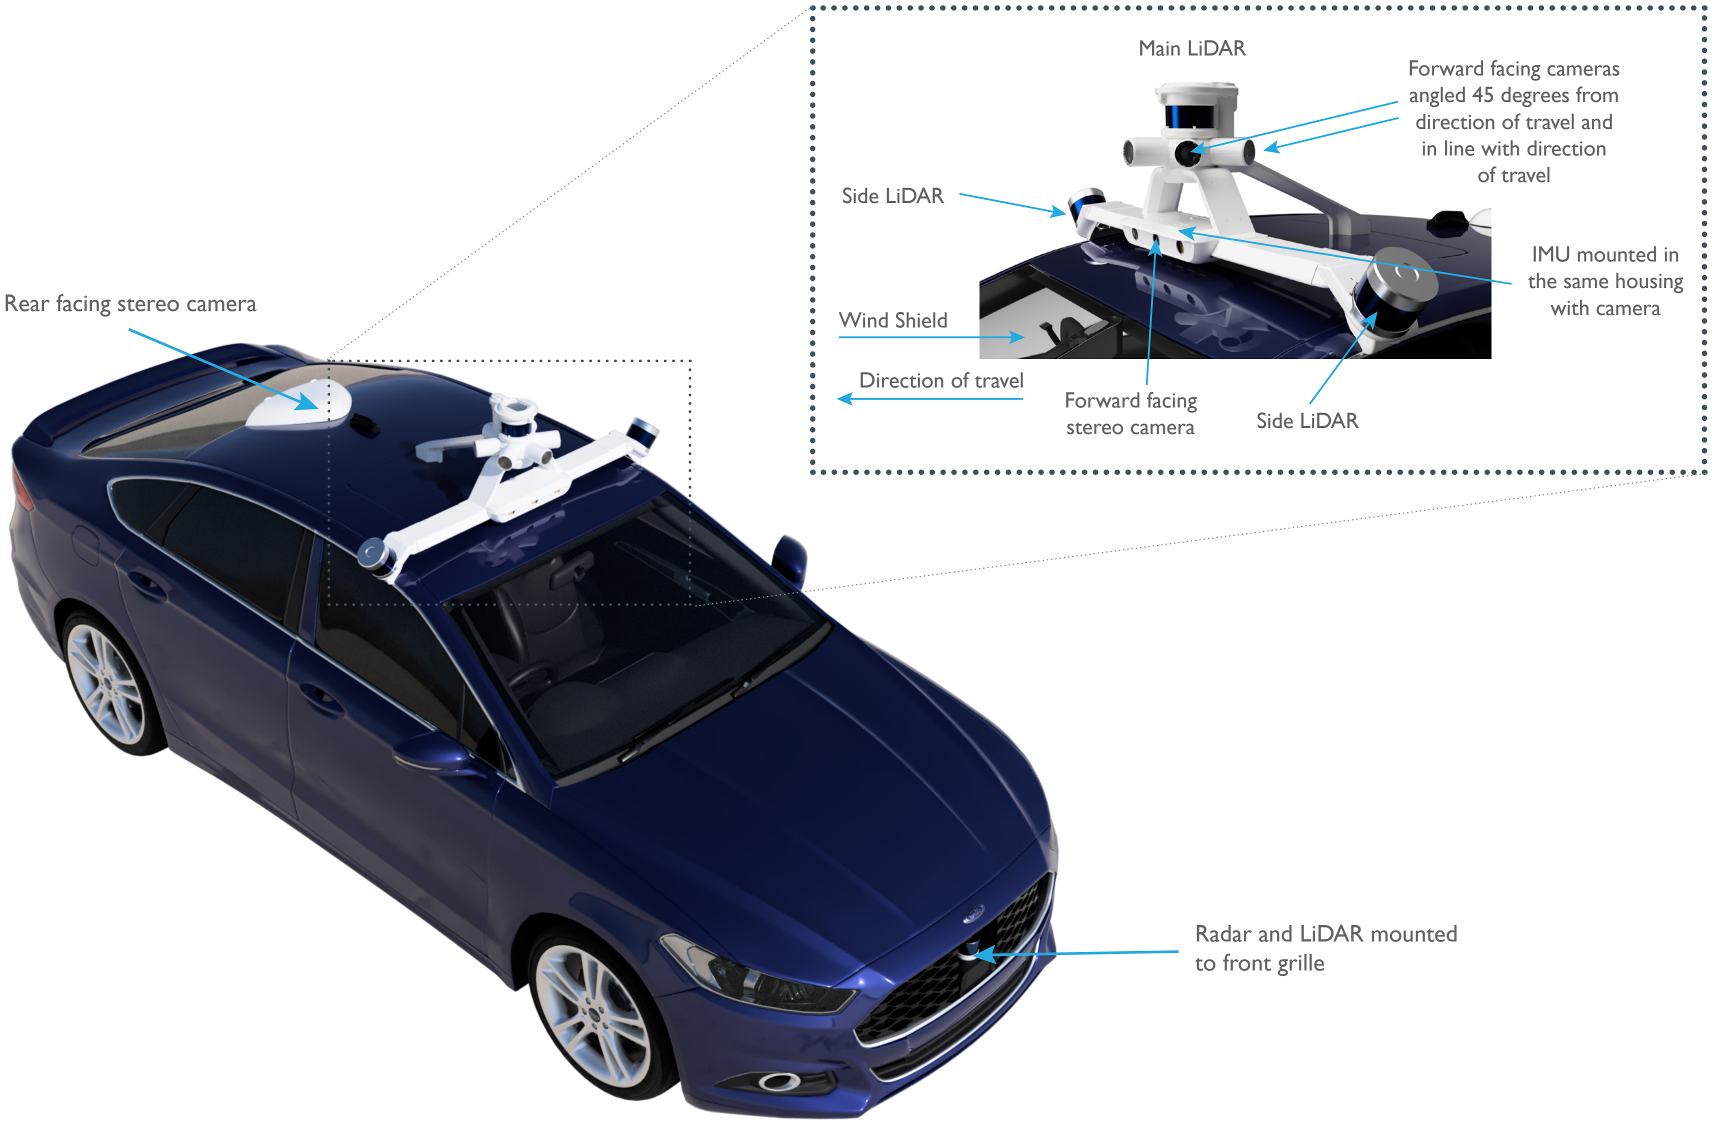

Conclusion

The SHIFT project, a collaboration between Oxbotica, Imperial College London and Transport for London (TfL), has achieved the objectives it set out: to understand and define the key considerations and impacts of deploying autonomous vehicles (AVs) safely and efficiently in urban environments.

These are laid out in the SHIFT Autonomous Vehicle Deployment Report, including a model to understand the potential impact of AV services on the London road network, a robust fleet management system capable of handling the required data, an AV operator training template and a build order for AVs.

Already, Oxbotica has utilised the SHIFT Autonomous Vehicle Deployment Report as part of Project Endeavour, a Government-backed mobility project designed to accelerate and scale the adoption of AVs services across the UK through advanced simulations alongside trials on public roads in three major cities. The build order and data infrastructure developed through SHIFT has been deployed on a fleet of six Ford Mondeo vehicles which are autonomously completing a nine-mile round trip from Oxford Parkway station to Oxford's main train station.

Similarly, Imperial College London and TfL are looking to continue to work on modelling opportunities. The findings from the Delos Urban Mobility model developed during SHIFT will be used to improve TfL's own multi-modal demand model; Model of Travel in LondON (MoTiON). Working with outputs from the Delos Urban Mobility model developed during SHIFT, gives the opportunity to define some of the parameters in MoTiON's AV functionality, for example passing data from MoTiON into the SHIFT model, means that the Delos Urban Mobility simulation will output some key data that TfL has previously been unable to generate.

MoTiON will now be able to include individual waiting times, journey times and detour time (for shared trips) as well as how many existing private hire vehicle demand trips would be served in a scenario where users are highly likely to share journeys. With this updated functionality in MoTiON, TfL can perform tests to understand the impact of these newly-specified AVs on travel demand in London.

SHIFT

THE AUTONOMOUS VEHICLE DEPLOYMENT REPORT