

WHEN WILL CARS DRIVE THEMSELVES?



The new Mercedes S 500 autonomous car uses sensors already fitted on the production S-Class to 'see' obstacles and interpret roadside information and a sophisticated route pilot to navigate the way. In September 2013 the pilot car took on traffic lights, roundabouts, pedestrians, cyclists and trams, as it navigated its way from Mannheim to Pforzheim, in western Germany

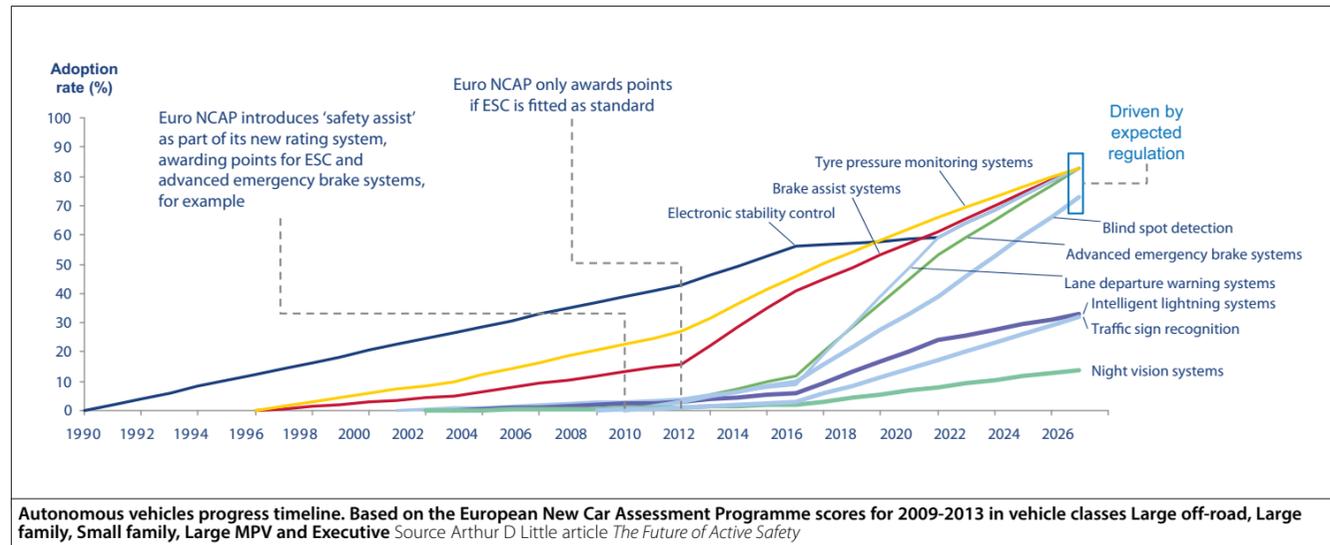
There are many claims made about the progress of autonomous vehicles and their imminent arrival on UK roads. Professor Tim Gordon, Head of Engineering at the University of Lincoln, and Mathias Lidberg, Associate Professor at Chalmers University of Technology, look behind the headline hype to see what progress has been made and how measures already implemented have increased automation.

Cars are getting smarter: bristling with actuators and sensors, they are increasingly sophisticated and computer-controlled. The new technologies include automatic collision avoidance by braking, automated parking (even into tiny parking spaces) as well as automated driving in traffic jams.

Several companies – including Google, Volvo, Tesla and Daimler – are demonstrating fully self-driving vehicles that use computers

rather than humans to do the driving. For safety reasons, these 'driverless' or 'autonomous' vehicles have test engineers closely supervise their progress. There are, however, intentions to make self-driving cars available to the general public. The US states of Nevada and California recently introduced legislation to allow testing of self-driving cars on their highways, and the UK government recently announced a similar initiative.

While there have been plenty of headlines about these autonomous vehicles, handing over complete control of our cars to computers is just one end of a spectrum of automation. More significantly, vehicle automation offers a growing portfolio of limited but useful capabilities. We can expect to see further evolutionary development, with cars becoming progressively smarter, and therefore providing more assistance to the driver behind the steering wheel.



BRAKING BREAKTHROUGHS

An early example of the technologies needed for self-driving vehicles is the antilock braking system (ABS). Developed in the 1980s, ABS technology monitors wheel speeds and controls brake pressures to prevent wheels from locking up under hard braking. ABS may not be considered particularly 'smart' these days, but its sensors and actuators provide a gateway towards more sophisticated systems. For example, in electronic stability control (ESC), a computer intervenes with the brakes to correct for loss of stability, which is not uncommon on slippery or icy roads. Compared with ABS, ESC interacts with the driver more conspicuously, correcting the way the car moves using 'feedback control'. ESC compares

actual with expected vehicle motions. Using error-correcting algorithms, it then prevents the car from spinning dangerously or sliding sideways. ESC requires further sensors to monitor the vehicle's motion, as well as sophisticated software, but the additional costs are low compared to the original installation of ABS.

Taken together, ESC and electronic engine control make it relatively simple, and affordable, to use engine and brakes to control a vehicle's speed. By adding forward-pointing radar, a vehicle can follow a driver-selected speed and adapt to traffic by slowing down when there is a slower vehicle in front. This first step towards automated driving emerged as adaptive cruise control (ACC) in the late 1990s.

A growing list of car manufacturers now offers ACC,

including Cadillac, Jaguar, Honda, BMW and Volvo. Fully tested and proven as useful and safe, it tends to reduce tailgating as well as cutting the workload for the driver.

BEYOND ACC

Over the past few years, collision avoidance and crash mitigation systems have come onto the market as further developments of automatic speed control. In these systems, the car, or lorry, automatically applies its brakes to avoid a crash or reduce the severity of an anticipated impact. This is another example of existing actuators and sensors from ABS and ACC being co-opted for a new purpose, with the vehicle's computer given authority to apply the brakes even when there is no command from the driver. The system's role is elevated from reducing the

driver's workload to preventing crashes. In fact, it is part of the evolving new generation of 'autonomous safety functions' in passenger cars and heavy vehicles.

Speed control is only half of the picture, however. For full self-driving, a vehicle also needs automated directional control. Directional control usually comes from turning the steering wheel. Typical demonstrations of self-driving show the steering wheel moving 'by itself'. As with speed automation, the key requirement is to sense the car's motion relative to the road, and/or relative to other road users, and then to apply path corrections to reduce or remove errors between the actual and intended paths.

Automated speed control is one component of the broader concept of drive-by-wire. This concept of automated car

control borrows its thinking from the fly-by-wire approach of modern aircraft. As in an aircraft, the technology handles the car's braking, steering and propulsion. A complete drive-by-wire system combines individual systems that, for example, steer-by-wire or throttle-by-wire – see *What's behind the wheel?*

Since it is usual – though not universal – to have a direct mechanical connection between the steering wheel and the direction of the front wheels, there is a design challenge to best assist the driver with steering. One option is to use an electric power assistance system (EPAS) to adjust the steering. This is already common in the power steering of more recent vehicles, where an electric motor replaces the hydraulics.

When used in a more interactive or 'intelligent' manner, the car will appear to 'want to steer' of its own accord. Clearly, with hands on the wheel, the driver will perceive the car as having its own intention. Provided drivers are sympathetic to that intention, they will likely allow the steering action. EPAS is already used in self-parking systems where the steering action is fully automated – drivers remove their hands and simply control brake and accelerator during parallel parking. On the other hand, should unexpected steering torque be experienced during regular driving, this could disturb the driver and cause them to resist, which could have dangerous consequences.

There are other options for the car to control its direction,

with different implications for the interaction with the driver. For example, with active front steer, an electric motor controls an offset between the steering wheel and the front road wheels. This corrects the steering on the road without creating a corresponding movement of the steering wheel.

In another approach to automated steering, 'steer-by-wire' removes the mechanical connection between the steering wheel and road wheels and so offers further scope to modify the interaction with the driver. This type of electronic control can be used to steer the rear wheels as well as the front, giving more freedom to designers. Yet another option is to create a left/right bias in the brakes or driveline; for example, by driving the right wheels

harder than the left, or braking the left wheels more than the right, the net effect is to turn the vehicle to the left.

All the approaches mentioned have their merits and implications in terms of cost, complexity, effect on the vehicle's motion and, perhaps most importantly, their synergy with the human driver.

SHARED CONTROL

Whatever the mechanism, sensors and actuators can provide an automated 'steering function'. As with speed control, the system can function in a self-steering mode only if there is a reference direction/path to follow. This direction can be determined either by the vehicle alone, acting in autonomous mode, or, working



WHAT'S BEHIND THE WHEEL?

Intelligent vehicles control themselves not through mechanical devices – by

turning a steering wheel or pressing brake pedals, for example – but by sending electrical signals that control the steering wheel and brakes through electric motors, via a

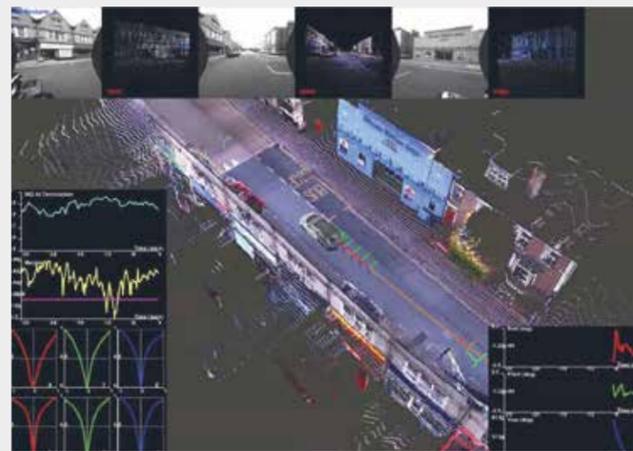
set of technologies collectively known as drive-by-wire. These electrical signals are in response to information from an array of sensors that detect the car's movement and direction, along with information about the local environment such as the presence of other vehicles, for example.

Nissan already has a steer-by-wire system in production. The Direct Adaptive Steering on its Infiniti Q50 Hybrid and Q50 3.7 models automatically makes small adjustments to the steering. For example, using what Nissan describes as a 'straight-line stability system', the steering can respond

automatically to crosswinds without disturbing the driver.

As a backup, if an electric motor fails or loses power, a clutch connects a mechanical linkage. This backup – having a redundant system that can take over – is important to ensure safe and fault-tolerant operation of the vehicle.

Redundant systems do not have to be mechanical. In a throttle-by-wire system, engine power is controlled purely via electrical signals. With redundant sensors and power supplies, there is no mechanical backup. This approach is already almost universal on modern cars.



The Oxford University RobotCar uses onboard cameras and laser scanners to navigate using a 3D map of the city.

WHERE AM I?

An intelligent autonomous vehicle needs to know where it is in the world. It needs to be able to fix its current map location, together with the positions of relevant objects such as road edges, cars, pedestrians and holes in the ground. At the moment, this is not possible: GPS satellite technology is unreliable in urban streets as buildings block the satellite signals, and in any case, standard GPS is not precise enough for controlling a vehicle.

Researchers in this field can use extra hardware to reduce positioning errors down to a few centimetres, but the costs are prohibitive for production vehicles. Sometimes an ad-hoc local 'map' can be created onboard the vehicle. For example, in lane-keeping control, cameras 'see' lane boundaries and provide the local road geometry, while relative positions of nearby objects

can be found from radar or laser ranging devices, often backed up with cameras performing image recognition.

A surveyed geographical map can provide additional information, especially detailed road geometry, the locations of traffic lights and the ability to plan journeys. However, to make use of surveyed maps the precise position of the vehicle is needed. There is currently no single accepted way to do this for economically-feasible intelligent production vehicles.

One approach, being developed in the UK by Professor Paul Newman at the University of Oxford, is to use low-cost cameras and laser scanners mounted on the vehicle to detect and record the locations of fixed objects on 'familiar' routes, so that precise positioning is possible in subsequent trips. The Oxford Mobile Robotics Group are combining positioning and map building into a single process – see <http://mrg.robots.ox.ac.uk>

in shared-control mode, as an adjustment to what the driver does with the steering wheel. In either case, the vehicle needs to know its position and motion relative to some map of the road and potentially to obstacles, including other vehicles – see *Where am I?*

In one form of shared control, lane-keeping assist or lane-keeping aid (LKA), now available in cars such as the Volvo V60, a camera determines the car's position relative to the lane markings. The electronic steering function then keeps the car in the lane. The steering control is gentle enough to allow the driver to override the system for lane changing or to avoid an obstacle, in much the same way that ACC will defer to the driver for braking or acceleration. In both cases, the responsibility for safety and legal liability stays with the driver.

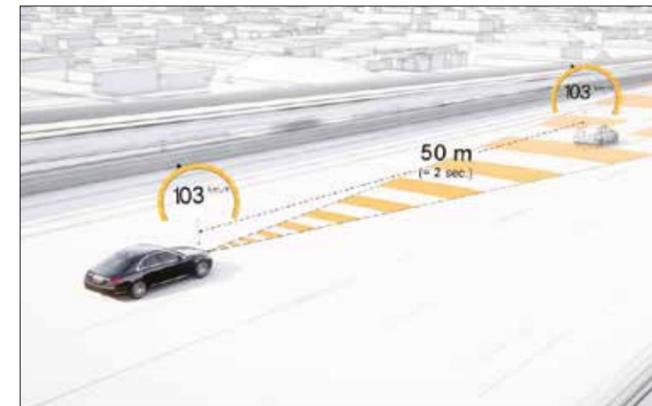
LKA may also monitor whether the driver is holding the steering wheel. If the driver removes their hands for more than a short time, the system will give a warning and then disconnect itself. Automated steering can also be a safety enhancement. For example, side-pointing radar can detect another vehicle in the blind spot in an adjacent lane so that the system will resist the driver's steering action if the driver starts to make an unsafe lane change. It is currently difficult to predict the actual safety benefits of such systems, even using so-called field operational tests.

We have seen how new 'higher' levels of automation build on existing technologies. Now there emerges a new and thought-provoking opportunity. ACC and LKA are both

commercially viable systems, available today, so switching them on together could immediately give us a self-driving car! However, both systems have limited capabilities: neither is designed to take full authority from the driver. It follows that we still require the driver to supervise: feet off the pedals, hands off the steering wheel, but eyes firmly on the road. And, importantly, the driver/supervisor should be mentally alert and primed and ready to take control at any instant.

Several companies are developing such 'semi-autonomous' systems with the motivation of further reducing the driver's workload. However, the shortfalls of boredom and distraction during a low-level supervision task are well known – which is why lifeguards at swimming pools are rotated after a short time of supervision.

When the driver has to remain vigilant (here we still say 'driver') one doubts there is much genuine workload reduction. On the positive side, the driver can relax their arms and let the steering wheel move by itself. On the negative side, there is the new task of supervising and monitoring the automated system, and deciding if and when to take back control. The gain is a reduced physical demand, while the new burden is largely cognitive. Crucially, the semi-autonomous vehicle does not alleviate the normal task of anticipating traffic hazards. Added to this are the dangers of overreliance on the system and being distracted by other tasks, potentially leading to new crash risks. This acts as an incentive for creating forever more ingenious



Autonomous driving in the Mercedes S-, E- and C-Class of 2013/2014. With the Mercedes-Benz intelligent drive Stop and Go pilot car, the vehicle moves autonomously at low speeds in traffic jams. It has longitudinal and lateral control including lane keeping in curves and lane mark detection. It is able to observe vehicles ahead using warn mode. With its intelligent hands-on detection, the driver is called back into the loop when needed

systems which reduce crash risk, while at the same time allowing 'eyes off the road' driving.

DRIVING SKILLS AND RESPONSIBILITIES

It is easy to underestimate how good we humans actually are at driving. According to US traffic statistics (Traffic Safety Facts, 2012), out of roughly three trillion vehicle miles travelled, there were a total of 5.6 million police-reported traffic accidents, 30,800 of them involving a fatality. This equates to roughly one crash per 500,000 miles and one fatal crash per 100 million miles. This compares to the currently quoted million-plus miles of accident-free (supervised) travel for Google's self-driving cars, where the development team selects the roads and timing, intervenes when necessary, and presumably does so with high attention levels.

Driving is usually thought of as having three levels: strategic, tactical and control. Strategic driving deals with choice of destination, route, timing and other forward planning – something that current mapping and navigation software can deal with very well. The control or skill level deals with stability, staying in the lane, slowing down for a curve, executing a lane change: the kind of things that are relatively easy for current self-driving technologies. The intermediate, tactical level deals with making decisions in traffic, deciding on which lane to be in, recognising and anticipating potential hazards and so on.

For human drivers, the tactical level requires experience and attention. It is also the level that is most demanding for self-driving cars. It is not feasible to replicate all aspects of human driving at the tactical level, at least not without artificially intelligent computers that are beyond anything we currently



In May 2012, Europe's first ever autonomous vehicle road train took to the road outside Barcelona, Spain. The convoy of autonomously driven Volvo vehicles followed behind a lead vehicle driven by a professional driver. The five vehicles drove on a public motorway among other road users at 85 km/h, with a gap between each vehicle of 6m. During the three-year project the vehicles drove a total of 11,000km © Ricardo

have. We might even formulate a 'Turing driving test': if a self-driving car behaves in all traffic conditions such that other road users cannot tell any difference from a human driver, then we may be inclined to assume we have a safe and acceptable system.

If we cannot yet put sufficient human-like driving intelligence into a box of electronics, there are two complementary ways forward. First, we can make best use of the available 'safety net' of sensors and actuators to avoid collisions and reduce the risk of crashes. Over time, using crash data and improved scientific analysis, we can evaluate and verify the safety benefits of such systems. Since the safety nets developed to protect human drivers can protect us during autonomous driving, the driving safety benefits can be extended similarly. Secondly, we can try to simplify or control the driving environment to accommodate these infant electronic driving systems.

There are examples of where this has been done already. For example, in August 1997, a large R&D project ended with a live proof-of-concept demonstration of a prototype automated highway system, with 20 self-driving cars operating over a four-day trial period on dedicated lanes of Interstate 15 in San Diego, California. With speed and directional control both automated, and permanent magnets were embedded in the road guiding the vehicles. The resulting 'slinky effect' or 'string instability' can be unnerving when the lead vehicle accelerates – following vehicles accelerate progressively harder to catch up – and could lead to collisions whenever the lead vehicle brakes suddenly. Interestingly, the 1997 demonstration showed that automated vehicles can control platooning more precisely than human drivers. With close spacing achieved at highway speeds, the demonstration showed potential advantages for



By mid-2017, it is planned that 100 fully autonomous cars will be operating in Milton Keynes, sharing pathways with pedestrians and equipped with sensors to avoid collisions. The LUTZ (Low-carbon Urban Transport Zone Pathfinder) driverless pods, which can carry two people each, will be able to travel at a maximum speed of 24km/h © Transport Systems Catapult

Pods. The system, akin to a flexible rail system, is a long way from the freedom of general purpose self-driving cars, but it demonstrates real value to customers and is already economically viable.

A different pod system is planned to ferry passengers between the rail station in Milton Keynes and nearby shopping centres. Fully autonomous pods will be trialled in Milton Keynes early in 2015 by the Transport Systems Catapult. Ultimately ferrying passengers between the railway station and nearby shopping centres, the pods will run at up to 24kph and be tested on a pavement route agreed with Milton Keynes Council to see how they interact with the public.

There is a clear distinction between autonomous safety functions, such as collision avoidance by automatic braking, which are expected to be always vigilant but rarely active, and autopilot functions such as ACC, which automate some or all of the driving task for extended periods. It is theoretically easier to assess the safety implications of the first set of features. We only need to evaluate their capabilities in certain well-defined situations: for example, when the vehicle in front brakes hard.

In the second set, which may operate for extended periods of time, it is important to evaluate everything that may go wrong, a much harder task. In both cases, the scientific basis for predicting the safety of new technologies lags the development of the technologies themselves. It is not useful to rely on counting crashes to evaluate the safety of self-driving cars.

With increasing levels of driving automation, it becomes clear that we also need better sensing, monitoring and regulation of the driving environment, particularly so that we can accommodate the limited tactical capabilities of automated driving systems. We need smart roads, not just smart cars. One option is to provide dedicated lanes for automated driving, but this is expensive and can hardly fulfil any 'restaurant and back' aspiration. More precise and widespread

increasing traffic flow and reducing aerodynamic drag among the following vehicles.

In a recent EU project, Safe Road Trains for the Environment (SARTRE), led by Ricardo UK, a heavy truck with a lorry driver at the wheel acting as platoon leader, was followed by suitably equipped and automated passenger vehicles. Again, with highway speeds and close spacing, this 'mother-duck/baby ducks' scenario allowed drivers of following vehicles to enjoy eyes-off-the-road time. The approach is broadly feasible on existing highways. It simplifies the driving environment for each 'baby duck', including at the level of tactical driving. While the energy savings may not be large, SARTRE is an example of using a

novel approach to deal with the complexity of the traffic environment.

A third and very different example of controlling the driving environment goes beyond the demonstrations of concept and further highlights the role of ensuring safety via a highly controlled infrastructure. Since 2011, personalised rapid transit pods have ferried passengers between Heathrow Airport's Terminal 5 and a nearby parking area. With speeds up to 25 mph, self-driving electric cars run on dedicated tracks, with boundary fences, concrete barriers that separate adjacent lanes, CCTV and other real-time monitoring, plus the safety net of an independent collision avoidance system within the

WHAT'S NEXT?

What can we learn from the driver assistance systems and autonomous driving developed and implemented so far? The first point is that there are many levels of self-driving functions or features. In any analysis, it is important to distinguish between these levels in regard to authority, situation awareness, capability and requirements from the human driver or supervisor.

Aside from the pod-style integrated systems, there is always some interaction with a driver, during operation or in switching between driving and non-driving modes. Properly managing this interaction will be vital to the success and safety of semi-autonomous systems

We need better sensing, monitoring and regulation of the driving environment, particularly so that we can accommodate the limited tactical capabilities of automated driving systems. We need smart roads, not just smart cars.

GOTHENBURG 2017



The first cars bought by customers as part of the Drive Me project are due to be on the roads in 2017, with the project running until 2020. It is expected that there will be 100 cars using 50 km of the city's roads in and around Gothenburg. These roads, typically commuter arteries, include highway driving and frequent stop-go traffic

The city of Gothenburg sees itself as the world's first arena for cars in everyday road conditions. The city is home to the 'Drive Me' project, which brings together Volvo Car Group, the Swedish Transport Administration and the City of Gothenburg.

The project, with the full name 'Drive Me - Self-driving cars for sustainable mobility', brings together legislators, transport authorities, a major city, a vehicle manufacturer and real customers. The project focuses on such areas as how autonomous vehicles bring societal and economic benefits by improving traffic efficiency, the traffic environment and road safety. The scheme is looking at the infrastructure requirements for autonomous driving, establishing those traffic situations suitable for autonomous vehicles and

checking how surrounding drivers interact with a self-driving car.

Drive Me began in April 2014, with the first tests of a modified Volvo Sedan S60 on public roads. The cars are equipped with a forward-looking camera and laser scanner mounted in a pod behind the rearview mirror; a radar system buried in the nose of the vehicle; and sonar emitters placed at strategic points around the car. They have GPS on the roof and tests have shown that the cars can handle lane following, speed adaptation and merging traffic by themselves.

The aim is that the vehicle's autopilot technology takes care of all driving functions so that the car itself handles all traffic scenarios, including leaving the traffic flow and finding a 'safe harbour' if, for any reason, the driver cannot regain control.

monitoring and regulation of the general road network will also be expensive, as well as being unwelcome where driving is regarded as a basic 'freedom'.

The history of smart technology on cars has been evolutionary, and seems likely to continue that way. The first full demonstration of self-driving highway vehicles took place 20 years before the next major effort – see Gothenburg 2017. There seems to be no reason to expect a rapid revolution any time soon. The most immediate of all the technical challenges is to find rigorous ways to prove the safety of self-driving systems in cars, especially in an unstructured driving environment. Safety will continue to be a significant

motivation for the technology, in the form of emerging autonomous safety functions in cars and heavy vehicles.

Ethical and legal questions may arise when cars take control and make decisions that affect the life and death of humans. Social, political and economic factors will strongly influence the pace of change. We should not be impatient: 30 years ago, a few engineers set out to improve braking performance on slippery roads and started connecting ABS hardware to simple computers. This started an evolutionary process that is leading towards sophisticated vehicle automation. Our cars are gradually turning into robots: it makes sense to take time and effort to get it right.

BIOGRAPHIES

Professor Tim Gordon is Head of Engineering at the University of Lincoln. He was at the University of Michigan, Ann Arbor, between 2003 and 2013, as Professor of Mechanical Engineering and Research Professor at the University of Michigan Transportation Research Institute. For six years, he chaired the committee on Vehicle-Highway Automation at the US National Academies Transportation Research Board, and is now co-chair of the IEEE Technical Committee

on Intelligent Vehicular Systems and Control. He is also Vice-President of the International Association for Vehicle System Dynamics.

Mathias Lidberg is Associate Professor at the Department of Applied Mechanics at Chalmers University of Technology, Sweden. He organised an international workshop on Vehicle Autonomy in August 2014, under the auspices of the International Association for Vehicle System Dynamics, involving leading academics and industry experts in the field.

The authors would like to thank Michael Kenward for his help in the writing of this article