BEHIND THE CURTAIN

When Driverless Car Demonstrations Are Less Than They Seem

AD PUNCTUM

October 2017

To The Point: A Summary Of This Report

This report was inspired by the many and varied demonstrations of self-driving vehicle technology over the past few years and the widening gulf between the appearance of capability and the reality. The aim is to inform non-specialists about some of the different methods used to enhance the apparent driving proficiency of prototype driverless vehicles.

Self-driving vehicles form an understanding of where they are and where they want to go using advanced versions of contemporary mapping and navigation systems -- mature technology. This includes dynamic route planning that changes course based on traffic conditions and road closures. Ideal paths derived from mapping are the foundation stone of nearly all (if not every) self-driving system. The disparity in capability between projects lies in how the car copes with differences between the ideal route and the actual environment. The best systems recognise objects and create an understanding of their real-time situation, together with predictions of how the scene might unfold. Lesser systems do not have this ability, or are capable only in simpler scenarios. This inadequacy can be disguised by the design of the demonstration (not that anyone would do such a thing). To explain the background clearly, this report covers the following areas:

A beginner's guide to object recognition -- a brief overview of what a self-driving artificial intelligence (AI) tries to identify in its surroundings and why.

An introduction to scene understanding and prediction -- an overview of how the artificial intelligence can use its understanding of the local environment to make driving decisions.

An overview of different demonstration events; relative difficulty and how to spot fakes -- four complexity levels:

- The parking lot demonstration
- The closed course demonstration
- The carefully selected on-road demonstration
- The high-confidence on-road Level 4 demonstration

This includes examples of how the demonstration can be simplified to make the vehicle appear more capable and some ways that you can investigate further. The issue is that, as shown in the table below, nearly all demonstrations appear sensational, so it is important to bring greater objectivity to the near certain euphoria felt on exiting the vehicle.

	Parking Lot	**Closed Course**	**Carefully Selected On-Road**	**High Confidence On-Road**
How impressive it is to non-specialists	Wow	OMG	#thefutureisnow	Shut up and take my money
Object detection required *	Limited	Limited	Good	Good
Object recognition required *	None	None	Some	Good
Scene understanding required *	None	None	None	Good
Ability to write a set of rules *	Easy	Easy	Quite Hard	Nearly Impossible
Ability to control demonstration conditions	Very Easy	Very Easy	Some	Very Hard
Competitive Level	Way behind	Way behind	Middle of the road	Near the front

* Minimum requirement to stage a convincing demonstration; no implication that past demonstrators have done this

The only conclusion is **buyer beware -- look carefully behind the curtain**. Very few people have travelled in a driverless vehicle and the experience remains impressive, even in circumstances where it is heavily staged. This report simply aims to assist objectivity in the face of thrilling and often seemingly compelling technology demonstrations.

A beginner's guide to object recognition

Self-driving vehicles make real-time decisions by combining the ideal intended path (normally derived from high definition mapping data) with information about the current state of the local environment. It must understand what the surroundings are and how that affects the ideal path to the destination. The first step is recognition and classification of the objects around the vehicle. This is done individually and then combined to build awareness of the overall environment, often called **scene understanding,** and **free space** (the amount of road that isn't blocked or in use by someone else.

Artificial intelligence (AI) systems are being trained to recognise different objects from multiple angles. Although opinion about the minimum level of recognition necessary is divided (e.g. is it important to know that an approaching vehicle is a Tesla Model S or simply that it is a car?), most AI systems are being trained to attribute object images to multi-layered hierarchies, as demonstrated in **Figure 1**.

FIGURE 1 | Say What You See

Source: Wikimedia Commons

This is an adult Labrador, which is a Labrador, which is a dog, which is an animal

Source: Wikimedia Commons

This is a 2015MY Ford Focus RS, which is a Ford Focus RS, which is a Ford Focus, which is a Ford and also a compact car, which is a type of car, which is a type of vehicle

A potential approach is for a vehicle to simply identify objects with virtually nil attempt to categorise them differently between, for instance, a car, a dustbin and a dog. In theory these are all simply objects to be avoided. Although there are some situations when this can help simplify the approach -- especially over longer ranges -- it is impractical to adopt this method for driving decisions regarding objects that are close by. There are two primary reasons for this:

Firstly, without understanding what the object is, it is not possible to make any prediction of how it is likely to act over the next few seconds or how it will react to the actions of the vehicle. A self-driving car would have to act as if every dustbin it passed was potentially a child about to run into the street. The result would be a driving style that fits the lowest common denominator of risk -- e.g. no faster than a walking pace and only making lane changes or turns when large gaps open up.

Secondly, machine vision works in a way that is quite different to human vision. Therefore, intuition around object recognition does not perfectly apply. Machine vision begins by looking at a single point, then seeing what surrounds that point and then seeing whether the point with its surroundings matches a pattern that has been observed before. This basic premise holds whether the sensor involved is camera, radar, lidar or ultrasonic -- all inherently output as points or pixels. For this reason, machine vision needs to have some level of object recognition because otherwise -- in an extreme case -- it would see the entire surroundings as a single large object. If AI had feelings, it would feel trapped forever.

FIGURE 2 | Who's A Clever AI Then?

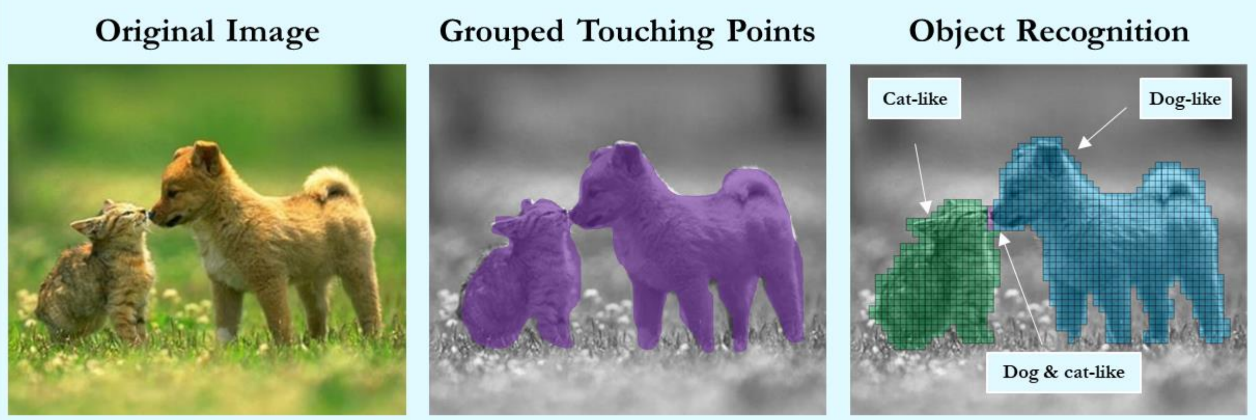

Source: Flickr / kitty.green66, Ad Punctum Analysis

Figure 2 uses a simple example to show how improved object recognition improves understanding. The first image is what humans see -- a cat and a dog, their noses touching. The second image is what a poorly trained AI would recognise.

It detects a group of points that are distinct from their surroundings but not that there are two separate objects. The AI has no way of knowing -- from this single frame -- that there are two independent entities with different behavioural characteristics. One will be your friend for life and the other will scratch your eyes out. The third image is a rough approximation of how machine vision works -- the software finds pixels that have dog-like properties (marked in blue) and are next to other pixels that are dog-like. It does the same for cat-like pixels (marked in green). There is a group which are both dog-like and cat-like (marked in purple where the noses touch). The AI can assign a high probability to this being a point where the two are meeting. In concept this is the same for pixels or point clouds.

Self-driving AI uses a continuous video feed rather than single images. This is in some ways easier because more frames provide more attempts at identification, but it can also be more difficult. The AI must reconcile the object looking different between frames (e.g. changes in lighting or because the object is moving) and because the vehicle is normally moving towards or away from the object.

Figure 3 lists some examples of why detailed object recognition can be useful to a self-driving vehicle. On the left-hand side is a hierarchy of different objects that might be encountered on a journey. Even before trying to gain the full understanding of what the overall scene is, object recognition informs the self-driving AI about probability of unexpected hazards and provides a comparison to the existing map it is using to determine the best route. The self-driving vehicle can immediately see which objects need to be monitored closely and which are more benign. In addition, having identified which objects are in the road, the vehicle can understand the free space available to it.

FIGURE 3 | I Spy With My Self-Driving AI…

				Could impact my ideal route?	Potentially cause unpredictable hazards?	New information versus my map?	Help confirm my position?
Vehicles	Car	Ford	Various	Yes	Yes	Yes	No
		VW	Polo	Yes	Yes	Yes	No
			Golf	Yes	Yes	Yes	No
			Various	Yes	Yes	Yes	No
	Van	Various	Various	Yes	Yes	Yes	No
	Heavy Goods Vehicle	Various	Various	Yes	Yes	Yes	No
	Bus	Various	Various	Yes	Yes	Yes	No
Pedestrians	Adult	Man	Man with a hat on	Yes	Yes	Yes	No
			Man with a pram	Yes	Yes	Yes	No
			Various	Yes	Yes	Yes	No
		Woman	Various	Yes	Yes	Yes	No
	Child	Boy	Various	Yes	Yes	Yes	No
		Girl	Various	Yes	Yes	Yes	No
	Animal	Various	Various	Yes	Yes	Yes	No
Road Sign	Temporary sign	Lane closure	Various	Potentially	Yes	Yes	Highly unlikely
		Narrow lanes	Various	Highly unlikely	Yes	Yes	Highly unlikely
	Permanent sign	Speed limit	40 mph	Almost zero	Highly unlikely	Highly unlikely	Yes
			55 mph	Almost zero	Highly unlikely	Highly unlikely	Yes
			Various	Almost zero	Highly unlikely	Unlikely	Yes
		Road hazard	Animals in road	Highly unlikely	Yes	Unlikely	Yes
			Steep gradient	Almost zero	Almost zero	Highly unlikely	Yes
			Various	Almost zero	Almost zero	Unlikely	Yes
		Parking conditions	No Stopping 10am - 6pm	Almost zero	Almost zero	Unlikely	Yes
			Residents parking only	Almost zero	Almost zero	Unlikely	Yes
			Various	Almost zero	Almost zero	Unlikely	Yes
Kerbside Objects	Dustbin	Stationary bin	Various	Almost zero	Almost zero	No	Yes
		Moveable bin	Various	Highly unlikely	Highly unlikely	No	No
	Telephone booth	Various	Various	Almost zero	Almost zero	No	Yes
	Streetlight	Various	Various	Almost zero	Almost zero	No	Yes
	Tree	Various	Various	Highly unlikely	Highly unlikely	No	Yes
	Fire hydrant	Various	Various	Highly unlikely	Unlikely	No	Yes
	Traffic lights	Various	Various	Highly unlikely	Yes	Yes	Yes
Road markings	Lane markings	Various	Various	Highly unlikely	Unlikely	Unlikely	Yes
	Pedestrian crossing	Various	Various	Highly unlikely	Yes	Highly unlikely	Highly unlikely
	Stop line	Various	Various	Highly unlikely	Yes	Highly unlikely	Highly unlikely
Objects in the road	Traffic island	Various	Various	Highly unlikely	Yes	Highly unlikely	Yes
	Roundabout	Various	Various	Highly unlikely	Yes	Highly unlikely	Yes
	Traffic cones	Various	Various	Potentially	Yes	Yes	No
	Kerbside	Various	Various	Highly unlikely	Highly unlikely	Highly unlikely	Yes
	Traffic lights	Various	Various	Highly unlikely	Yes	Highly unlikely	Yes
	Lane divider	Various	Various	Highly unlikely	Highly unlikely	Highly unlikely	Highly unlikely
	Pothole	Various	Various	Potentially	Yes	Yes	No

Source: Ad Punctum Analysis

The difference between limited and advanced recognition will only show in more complex test scenarios. In simpler tests where possible encounters are tightly controlled, less advanced systems can give the appearance of real-world capability.

An introduction to scene understanding and prediction

Self-driving vehicles use their knowledge of their position in space and time (normally from a map) and recognition of the objects around them to make second-by-second decisions according to the following high-level principles:

- The ideal route according to mapping, traffic and other data
- The free space on the road that is available to the vehicle
- The current actions of other road users
- The potential actions of other road users

To apply this logic properly, the AI must therefore comprehend the scene it is currently a part of. It can then apply a predictive framework to the actions that they might take next. It is not strictly necessary to understand the scene to generate the ideal route, understand the current free space or track the current behaviour of other road users. A self-driving vehicle that does not have scene understanding can therefore appear very competent in a simple test because all the challenges it must overcome can be covered with a list of rules.

In theory all driving could be covered by an exhaustive set of rules and if successful then the need for scene understanding would reduce. In practice either nobody is pursuing this, or they have not made their intentions public. The problem with that approach is generating the full list of situations and then writing, coding and testing the driving rules to identify and remove any contradictions. It seems sensible to assume AI will need predictive abilities.

Figure 4 takes a subset of the object examples from Figure 3 and applies a very simple predictive framework.

FIGURE 4 | Guessing Game

		High Probability Action		**Potential Action**		**Information Known / Inferences Made**			
		Movement	Speed	Movement	Speed	...about this type of object	...about this specific object to help my prediction	...about the surroundings that could affect object's actions	Humans inside?
Ford	Various	Stay in lane	Fast	Many	Very fast	Performance range	Recent speed & direction	Near a turn? Path blocked?	Yes
VW	Golf	Stay in lane	Fast	Many	Very fast	Performance spec	Recent speed & direction	Near a turn? Path blocked?	Yes
	Polo	Stay in lane	Fast	Many	Very fast	Performance spec	Recent speed & direction	Near a turn? Path blocked?	Yes
Man	Man with a hat on	Stay on pavement	Slow	Many	Quite fast	Likely top speed & acceleration	Recent speed & direction	Near a crossing? Path blocked?	Yes
	Man with a pram	Stay on pavement	Slow	Many	Quite fast	Likely top speed & acceleration	Recent speed & direction	Near a crossing? Path blocked?	Yes -- Many
Woman	Various	Stay on pavement	Slow	Many	Quite fast	Likely top speed & acceleration	Recent speed & direction	Near a crossing? Path blocked?	Yes
Boy	Various	Stay on pavement	Slow	Many	Quite fast	Likely top speed & acceleration	Recent speed & direction	Near a crossing? Path blocked?	Yes
Girl	Various	Stay on pavement	Slow	Many	Quite fast	Likely top speed & acceleration	Recent speed & direction	Near a crossing? Path blocked?	Yes
Stationary bin	Various	None	Zero	Fall into road	Slow	Highly unlikely to move	Recent wobbles	Weather conditions	Highly unlikely
Moveable bin	Various	None	Zero	Roll into road	Slow	Highly unlikely to move	Recent wobbles	Just emptied?	Highly unlikely
Tree	Various	None	Zero	Fall into road	Fast	Highly unlikely to move	Recent wobbles	Weather conditions	Highly unlikely

Source: Ad Punctum Analysis

A framework such as the one above enables self-driving AI to follow an intuitive logic path:

- The map takes me along the high street, there should be two lanes available -- neither is theoretically preferable
- There are vehicles parked at intervals in the left hand lane -- I will stay in the right-hand lane and avoid obstacles
- The vehicle ahead of me is travelling at just under the speed limit -- I will match their speed and maintain a safe distance in case they want to stop
- Stay alert to risks -- the car behind me is tailgating me and just performed an aggressive overtake on the car that it was behind previously; there are two pedestrians standing at the side of the road 15 metres ahead of my current position

In this simple example, the logic flow concentrates on the information that the AI can determine for itself. If vehicles are connected to each other, the infrastructure and even the cloud it is possible to be far more sophisticated. For instance:

- The navigation system of the car ahead can tell the AV that it intends to go straight for the next 500 metres
- The pedestrian's calendar knows that they have a meeting starting in 90 seconds at a location three streets across from here and so might alert traffic nearby (in a way that protects their privacy) that they might rush across the road
- The infrastructure could notify vehicles that someone pressed the pedestrian crossing button 30 seconds ago and so could be growing impatient

Such information would enable a self-driving AI to attribute more accurate probabilities to the chance of the vehicle ahead suddenly changing position (unlikely unless in response to external stimulus), and pedestrians running across the road (both cases increase the likelihood).

The explanation above can only be used illustratively. Whilst this is in concept how a self-driving AI approaches a situation, it might not ever be possible to discover the set of rules the AI uses to drive (especially in a form that is man-readable). The reason is that deep learning creates a set of rules determined by the AI itself from training data, rather than written by humans. The positive of this is that the humans are relieved of coding. The downside is that since the AI arrives at these conclusions based on its training set, differences in the examples provided could influence the driving program. Two vehicles with identical deep learning software and sensor hardware could end up with a different driving style to one another, simply based on the training examples that they encountered. For this reason, the prevailing approach is to have a central AI deep learning program that downloads a copy of itself to a fleet of vehicles, rather than each vehicle having its own unique personality.

Given the risks above, it is important that self-driving development teams can demonstrate the methods they use to bring scale to their learning. They need techniques that allow multiple vehicles and simulations to create and test at the same time, but there must also be rigorous quality control in the training data to prevent the AI learning the wrong thing. In addition to the exciting bit of driving around to collect situation data, they need a properly catalogued and labelled library or images and sensor traces. Without this, the approach boils down to telling the AI to train itself not to crash. This is fine in theory but quickly creates a lowest common denominator approach (limp mode) in any challenging situation.

Figure 5 shows how a properly trained AI might understand a situation. The vehicle is travelling down a road and has identified the free space available to it (in green). This encompasses all the road going straight ahead plus the left-hand turn and some of the right hand turning. Ahead of it is a car which has been travelling in the same direction (in blue) and appears as though it intends to turn right. There is a column of oncoming traffic comprising three cars (in purple) and one lorry (in yellow). Their lane position indicates that they intend to go straight on. There are no pedestrians visible.

FIGURE 5 | Before And After

Source: Wikimedia Commons (Geograph), Ad Punctum Analysis

Having understood the scene, the AI can decide how to proceed. This is an easy example; the vehicle can continue following the map route with little risk of collision in the next few seconds. In this situation it might be impossible to tell whether the vehicle operated on a simple set of rules with no scene understanding, had a limited ability to discern the setting or created very good environmental awareness. Simpler tests will not expose this difference.

An overview of different demonstration events relative difficulty; and how to spot fakes

Currently, no one has performed, or publicly aspired to perform, a true fully autonomous drive test. This would involve long distance journeys that traverse urban and rural areas with a variety of road types and weather conditions. Coast-to-coast trips exclusively using highways where the video is only released once it went off without a hitch clearly do not meet such a standard. All demonstrations so far have been of vehicles with some degree of geo-fencing. This restricts the vehicle from going "out of bounds". There isn't any inherent wrongdoing here but it is an important point to understand because a geo-fenced vehicle will have more restrictive use cases than a vehicle which can truly go anywhere. Whilst there are plenty of profitable applications for vehicles with limited use cases, it has impacts on product design and lifetime value and it does not necessarily follow that a geo-fenced SAE Level 4 capable driverless vehicle will scale to SAE Level 5. This is because confining the training data to limited areas reduces the time it takes an AI to become fully competent (good) but may result in the AI creating rules that are simply unworkable in other territories (bad) and it may not be possible to get the AI to drive differently without re-learning all the new situations from scratch (maybe not bad, but time-consuming).

This report details the four broad types of demonstration seen to date.

Source: Wikimedia Commons, Flickr

Since in all situations, the self-driving vehicle's journey is based on following a detailed map (in essence, replaying a recorded trip), it is only in harder situations that the AI's competence in image recognition and scene understanding is properly tested. It is quite possible to perform simpler demonstrations in a way that impress the audience but prove little in terms of technology that is needed for more complex environments.

The self-driving anarchist's store cupboard

Each type of demonstration covered contains tips on how to test the ability of the technology on display, using a simple array of props.

Here is the equipment list to take with you (if you can carry it) to a self-driving demonstration. If you are a guest then using any element of it would be the height of rudeness. If you are going to be expected to get out your chequebook at the end of the demonstration, then introducing your own modifications to a test should be fair game.

- Extra traffic cones -- for changing the route to see how the car reacts, or creating impromptu stationary obstacles
- White and yellow duct tape (& scissors) -- good for making new lane markings
- Mannequins with clothes -- helps you to see whether object recognition is based on movement or real understanding
- Beach balls -- allows you to introduce moving obstacles that don't create damage
- A skateboard -- put a mannequin on it, roll it into the road and see what happens next! Plus, good for performing tricks on during any pauses in the demonstration and generally looking cool
- A bucket of dirty water -- sophisticated sensor sets hate it
- A very sturdy wooden section with ramps at either side -- in practice probably nobody will do this but it gives you a chance to see how the vehicle reacts to road furniture (does it stop because it thinks it is a body, drive over it at high speed or does the vehicle slow and drive over it carefully as you might expect?)

THE PARKING LOT

What happens

The car drives on a private piece of open land, often a car park with a road layout determined by cones. At various points on the course, people walk into, or near, the path of the car and it slows or stops, thus proving that it is completely autonomous.

How it is done

- The car drives on a pre-programmed loop -- mainly relying on mapping and other position data (possibly nothing more than a recording of a human driver on the same route)
- The car relies on simple systems that detect obstacles, but not what they are and then slows or stops entirely until after they have passed -- the object recognition isn't important because the vehicle has been programmed according to the demonstration tests it will perform
- The car uses very simple rules to decide how to proceed -- normally something along the lines of "if within X angle field of view then slow to a stop, if outside angle X but within Y field of view then slow down to no more than Z miles per hour and proceed until out of field of view"

Developers may push this even further by telling the vehicle that it will only be subject to dangerous interventions in certain areas of the course -- for instance to impress with driving prowess on a technically tricky part and then show a pedestrian on a more boring part. They might either use an area with no road markings to avoid confusing the car's lane recognition systems or instead remove any lane recognition programming that could confuse the decision making. Finally, the lanes on the course (whether signified by road markings or cones) will often be much wider than those you would expect on the road.

How to spot it

- The path that the vehicle takes cannot be varied -- you aren't allowed to move the cones to create corners where there was previously a straight
- The vehicle drives each loop with pinpoint accuracy on the route, except when it avoids obstacles, but something feels odd; it has small steering movements that a human driver might make but a computer would not
- Whenever it encounters an obstacle, it returns to the ideal path from the prior loops in an identical fashion (e.g. pull out from obstacle, drive 5 metres, then return to ideal path)
- The team always present the same type of obstacles (and something always the same colour; blue car, man in a blue jumper) -- if you roll a shopping cart in front of it, the car either doesn't react or simply freezes
- The car doesn't realise that it is on a parking lot rather than an open road -- developers can't explain how come it is happy to ignore all lane markings and follow the cones instead (or can't show "it's okay, I'm in a parking lot" coding)
- If you sit in the driver's seat and grab the wheel then the whole car either shuts down or ignores you
- The program always starts from the same position on the loop

What you should do next

If this is something that has been put together in a hackathon by a small team then be impressed. If this is something that a group have worked on for a long time and believe justifies significant investment in the next stage of their development then proceed with extreme caution. There are over a hundred companies who are quite literally miles ahead of this point. A team demonstrating this technology now has little hope of catching up and doesn't appear to have thought of an innovative way to leapfrog the competition.

THE CLOSED COURSE

What happens

The car drives on what looks like a conventional road. However it is on private land such as an industrial estate or a test course, rather than public roads.

How it is done

- The car drives on a pre-programmed loop -- mainly relying on mapping and other position data
- The car uses simple obstacle detection -- it may group objects in a very simple way, for instance: large objects that are moving fast; large objects that are moving slowly; large objects that are stationary where the map says none should be.
- The car uses simple rules to decide how to proceed -- normally something along the lines of "if approaching known crossroads then decelerate to a complete stop. Do not continue on route until at least Y metres of free space appears"

Using a closed course also makes it possible to program in the different sections to give an appearance of "go anywhere" capability and dynamic route planning. The other benefit is that road surfaces may be better maintained than public roads (e.g. easier to see road markings, no potholes) and that crossing or turning decisions are made simply on whether there is an object on a possible collision course (no need to take pesky traffic lights into account).

How to spot it

- The vehicle has pinpoint accuracy journey after journey, except for where it detects obstacles, and makes slightly odd movements of the steering from time to time as if replaying a human driver's route
- Although the vehicle has a visibly impressive sensor set it approaches all stop signs / intersections with identical deceleration, regardless of other traffic
- The team always present the same obstacles to the vehicle (and at roughly the same point on the route) -- and normally one at a time
- Even though it is on a private estate, the team have closed off the road to other users
- Although there are road markings, the vehicle doesn't appear to use them -- if you create new markings (e.g. a new stop line) with duct tape, the vehicle ignores it and carries on
- If you sit in the driver's seat and grab the wheel then the whole car either shuts down or ignores you
- The program always starts from the same position on the loop
- Although the test demonstrates a reaction to moving pedestrians, if you put fully clothed mannequins on the kerbside, it doesn't recognise them (indicating that it relies on movement for object classification)
- The test doesn't involve real world manoeuvres such as another driver cutting in or needing to overtake a slow moving but not stationary vehicle
- The vehicle behaves oddly when pulling out from stationary vehicles or obstacles -- for instances, whether the object is a car or traffic cone it waits 10 metres before pulling back in, a row of ten cones one metre apart confuses it

What you should do next

This stage of demonstration is not as big a step on from the parking lot as it may seem. The road conditions are known so all combinations of the ideal route are relatively straightforward to capture. The correct behaviour towards other vehicles and pedestrians can be governed with simple rules because interactions are relatively rare and the vehicle has lots of time and space to make the appropriate reaction. Being able to drive with rules based reaction is far easier than needing to introduce artificial intelligence in decision making and cope with multiple targets at once. Lots of teams are more capable than this one.

In theory, if the sensor set and AI are well sorted and have been rigorously tested on a closed course then real world driving in simple conditions should be nearly within the team's grasp and they should have reams of data comparing a real driver trace with how the artificial intelligence would have reacted in the same conditions.

Lots of questions need to be asked -- what have the team really learned since their parking lot demonstration? What are their plans for on-road driving? Which locations are good and bad for learning and why? What have they identified as particularly tough conditions?

MY BACKYARD

What happens

The car drives on a public road using a safety driver. The route probably begins at the company base.

How it is done

- The car drives on a pre-programmed loop -- mainly relying on mapping and other position data
- The car relies on simple systems to detect obstacles -- it may group objects in a very simple way, for instance: large objects that are moving fast; large objects that are moving slowly; large objects that are stationary where the map says none should be.
- The car uses very simple rules to decide how to proceed -- normally something along the lines of "if approaching known crossroads then decelerate to a complete stop. Do not continue on route until at least Y metres of free space appears"

Operating on public roads is a significant step but there is still room for companies to make circumstances easier. Wide, multi-lane roads, avoiding roundabouts, rare encounters with pedestrians and other traffic on tricky roads, traffic lights only on low-speed roads. It is possible to design routes that appear to cover a range of conditions but don't get anywhere near the worst case of each.

How to spot it

- Roads that are not multi-lane are fairly clear of pedestrians and other traffic -- cars parked on single lane roads are never encountered
- The journey starts and ends at the company base -- vehicle may not be able to pull into kerbside and stop
- The team doesn't vary the route, even by just a couple of streets, or new routes take a long time to program in
- The team won't go near a school at going home time
- The route is all one type of road (e.g. low speed, wide roads, few traffic lights)
- The team won't demonstrate the vehicle in the wet or at night
- If you walk in the roadside near the vehicle it doesn't appear to change its path (or the team forbid it)
- Journey time is less than 15 minutes
- The team appear very selective about who the safety driver is and the driver isn't allowed to answer any detailed questions -- even ones that the ride-along VP of marketing couldn't possibly know such as "when have you personally been the most impressed with how this vehicle handled a difficult situation?"
- The vehicle seems to have a substantial support team and be kept in an environment that resembles a clean room
- Will they let you throw your bucket of dirty water onto the sensors and then do another demonstration?
- Will they let you throw a bucket of clean water onto the sensors and then do another demonstration?

What you should do next

This team is solidly in the pack of more competent groups, but is it one that will be making a breakthrough? One approach to test this is to choose a different demonstration route that incorporates more difficult conditions and give the team a limited time to learn the conditions. The challenge can be enhanced by choosing variable weather conditions.

If the team is well funded, it should be now running a decent sized fleet and augmenting the real-world learning with some degree of virtual testing. The team should have a clear test program and be able to explain the degree of progress made and expected on specific challenges. They should also be confident enough to speak about several dead ends that they have discovered during the development process.

YOUR BACKYARD

What happens

The car drives on a public road using a safety driver. The team appears confident enough to demonstrate in a number of conditions.

How it is done

- The vehicle has a large geography programmed into it
- The car can determine a preferred path for each route and then adjust it to the scene understanding
- The vehicle can cope with complex traffic conditions and multiple targets at once; although the intervention of the safety driver is sometimes necessary, it is sometimes rare

This is the current state of the art in demonstration. With or without public passengers, the vehicle can navigate a substantial area (normally of several square miles). Since at face value the “my backyard” and “your backyard” demonstrations are identical, the points below are aimed at proving validity rather than simplicity.

How to spot it

- The team let you choose the start and end points and the route
- The destination can be changed during the journey
- The team are happy to let the vehicle drive past a school at going home time
- The car can drive down a street where traffic travels both ways sharing the same lane -- it can discern right of way but also defend its own right to priority
- The car can cover highway and urban driving in a single journey
- The vehicle can carry out difficult merge manoeuvres at various road intersections
- Performance is equally impressive at night on a revised course
- They let you throw your bucket of dirty water onto the sensors and then do another demonstration

What you should do next

You have just been for a journey in a state of the art vehicle. The imagination can now run wild with dreams of chauffeurs for all within a limited time and dollar signs from the first to market monopoly that is sure to emerge.

But be careful. This vehicle may still be a long way from commercially acceptable capability in the hands of the public. In addition, artificial intelligence that can traverse the local environment may have little in common with a go-anywhere, do-anything driverless vehicle. As impressive as the vehicle is, it is not a given that the technology you have just seen demonstrated will scale.

Five questions need to be answered:

1. Is this team on a path towards the ability to drive anywhere? It may be that they are not and this could be okay, but if they believe themselves to be then they must be able to explain the issues and timescales involved.
2. Is this team developing the technologies that can be used capably and commercially in the specific use cases that are most likely to emerge several years ahead of the more capable vehicles that capture the public imagination? Is there an appropriate cabin for passengers or goods to work in? Have operational issues of a fleet been considered?
3. How scalable is the artificial intelligence and sensor set under development? Is the equipment robust enough that it can be let out of the hands of trained specialists for more than a few hours? How long does it take for a new environment to be learned?
4. What level of system redundancy have the team developed? To operate safely, a driverless vehicle will need several independent methods of making decisions
5. How flexible is the technological approach the team are taking? At present, self-driving solutions have little in the way of commoditisation or commonality, yet history indicates that bespoke technological solutions soon wither in the face of mass consumerism. The rate and manner of this is uncertain so teams must have flexibility in their system design and indeed would benefit from actively planning for all elements to become obsolete at some point or other

In Closing: A Summary Of This Report

Buyer beware -- look carefully behind the curtain. Very few people have travelled in a driverless vehicle and the experience remains impressive, even in circumstances where it is heavily staged.

Self-driving technology starts with an understanding of where a vehicle is and where it wants to go. The fundamental technology behind this has been demonstrated over many years by mapping and navigation systems. These now incorporate dynamic route planning that can instruct a vehicle to change course based on traffic conditions and road closures. Provided that unplanned obstacles are not introduced, a vehicle can drive a route relying on mapping data alone and do an impressive job of steering, accelerating and braking.

Object recognition has an important role. A driverless car needs to be able to recognise quite a lot of detail in order to properly inform itself about the surroundings. Some of this is easy to fake or do without in less challenging tests but cannot be avoided in more difficult driving environments, therefore a simpler demonstration may provide no basis for a more advanced trial if shortcuts have been taken.

Complex decisions rely on scene understanding and prediction. Most of the driving that an autonomous vehicle undertakes can be covered with simple rules. The car is told to always keep a distance of X metres from the vehicle in front, always give pedestrians a gap of Y metres and pass them at a maximum of Z miles per hour. The list is long, but straightforward and can be established quickly. This is what makes highly automated driving possible in production vehicles today. The issue is in dealing with circumstances that have not been prescribed, or with so many other actors in the scene that the rules start to overlap (e.g. it becomes impossible to chart a course past the cyclist without being unacceptably close to traffic passing on the other side). To overcome this, the vehicle must understand a scene understanding so it can predict the actions of the various other vehicles, people and objects and decide a best course of action (which could of course still be to stop and do nothing).

Simple demonstrations are relatively easy to stage and don't necessarily represent breakthrough technology. Many people are still wowed the first time that they encounter automatic parking. A vehicle that uses maps to drive an entire journey can create an experience that seems orders of magnitude further into the future by relying on the same fundamental technologies. A rules-based approach can cover a wide range of everyday conditions and provide a very competent looking experience providing that the vehicle doesn't have too many things thrown at it (physically or metaphorically).

About The Author

Thomas Ridge is a managing director in London. You may contact him by email at tridge@adpunctum.co.uk.

For Further Contact

If you would like to discuss this report, please contact the author.

Credits

Cover Image: freshidea / Adobe Stock

AD PUNCTUM

To find our latest research, please visit www.adpunctum.co.uk

Follow @Ad_Punctum on Twitter

Publication Number: 17J-BRA-P000031