The Received Wisdom

Episode 27: How could self-driving cars change the world? - Part 2

Jack Stilgoe: This is the second of a pair of episodes of The Received Wisdom that take a deep dive into a particular technology – my recent obsession – Self-driving vehicles.

In the previous episode, we focussed on Phoenix, Arizona, which has become a laboratory for the testing of the technology. That city, with its car-friendly streets and sunny climate, has seen an explosion in the number of vehicles shuttling about. But it has also witnessed the tragedy of the first bystander to have been killed by a self-driving test vehicle.

In this episode, I investigate what the technology might mean for London. I speak to some amazing people – sociologists, historians, artificial intelligence pioneers and transport thinkers. We begin by discussing roads as *places* where our social lives play out.

–

Intro: Steve Gooding: I sometimes felt that the government's enthusiasm for making this country a testbed for driverless technology would give us all the benefits that Christmas Island got from being a testbed for nuclear weapons. You get all the pain, but you don't necessarily get much of the gain.

–

Jack Stilgoe: I'm standing by a Zebra crossing in Russell square, around the corner from my office at University College London. The story goes that this is where, in 1933, the physicist Leo Szilard was crossing the road when he got the flash of insight that a chain reaction could cause a nuclear explosion.

It seems like a good place to think about what technologies can do to societies.

A zebra crossing is a little bit of shared space carved out of roads that are otherwise clearly the dominion of the motor car.

As a pedestrian, I have the Right of way. The highway code tells drivers, quote, “you MUST give way when a pedestrian has moved onto a crossing”

That’s the rule, but it’s not quite that simple in practice. Crossing the road is often a form of negotiation. We rely on a shared, common sense, and we may make eye contact to make doubly sure.

The rules are not always followed, and they are interpreted very differently. This is most obvious if you travel to different places with the same rules but very different norms. In Italy, for example, visitors are taught not to expect cars to stop.

If you’re an engineer trying to develop a self-driving car, the zebra crossing presents a fascinating and difficult challenge. If a self-driving car detects a pedestrian at the side of a zebra crossing, how does it know whether that person is trying to cross? How long should it wait before continuing on its way? Should there be some form of indicators to show nervous pedestrians they have been detected and inviting them to cross? Maybe it would be easier for everyone if crossings had traffic lights rather than the ambiguities of some white stripes?

These engineering choices are, deep down, political choices. When new technologies arrive, succeed and grow in the world, the world always adapts. People have to meet the technology halfway. That’s what’s happened with other forms of transport. If it happens with self-driving cars, it could reshape our cities. This could be exciting, or worrying, or both.

In the last episode, we heard from people in Phoenix, Arizona. Ground zero for self-driving cars. In this episode, we’re back in London asking whether the technology fits this older, more complicated, or whether the city could be remade to fit the technology.

I spoke to Joe Moran, a historian of everyday life, who wrote a lovely book on roads called, appropriately, On Roads. I asked him what he finds so interesting about the road as a place:

–

Joe Moran: Well, first of all, it's inevitably a shared space… in the early days of the motorcar, the motor car was the intruder, and the people driving them were kind of privileged and in the minority, so it was all about protecting the pedestrian and protecting other road users. And then I guess, from about the 1930s onwards, and 1930s, quite a key era, because that's when you got a lot of the things that we recognised the sort of furniture of the road, like the pedestrian crossings and road markings and things like that. Partly because the road accidents were just accelerating in number and that time, so people were worried about that, but also, the motorcar was becoming more common so that it then became about policing the pedestrian, and, and kind of trying to teach them how to behave.

It's quite an anthropologically interesting thing - an indicator of cultural differences - because there are obviously countries where jaywalking is very frowned upon. But it's always been based on the kind of slightly self-flattering English idea of, you know, compromise and common sense that we can all get along and we just need to sort of, we just need to sort of rub along together.

Jack Stilgoe: Can I ask you about speed cameras, because it seems just such an interesting, it's an interesting moment where technology meets these rules and the cultures that we all agree that there are speed limits, and most of us would agree that there need to be speed limits. But as soon as you have a machine enforcing those speed limits, it makes even some otherwise quite reasonable people extremely angry. I mean, what's going on? Why do speed cameras annoy people so much?

Joe Moran: I think in a way, speed cameras, were just another episode. People were getting just as angry about things like seatbelts and, and breathalysers in the 1960s. But possibly, there is something about being mechanically watched by, you know, a sort of surveillance state. But I think there's always been a thing about speed as well speed limits as being not fixed. That there's a there's a kind of give and take either way, because that was a that was a lot of the objection to speed cameras”

–

Jack Stilgoe: Pedestrian crossings as one place where these things get worked out

–

Joe Moran: I suppose the problem with a lot of the crossings is that, again, there's a sort of grey area. It's when the when the pedestrian steps on to the crossing, pedestrians

tend to be quite understandably nervous because a car can do a lot more damage than a pedestrian. From about the 1930s to the late 60s when they brought in Pelican crossing, they were trying to tweak that. It's a sort of English compromise, I suppose. It fits in with this sort of newer ideas about shared space, which came from Holland, about getting rid of traffic lights, and pedestrian crossing, things like that, but actually just kind of giving a sense from the architecture of the road and how it narrows and what it looks like making pedestrians cyclists, motorists, make it seem as though this is a space where you have to slow down.

–

Jack Stilgoe: When the highway code was first published in 1931 - just in time to tell Leo Szilard how to cross the road, the transport minister of the time called it a ‘code of good manners’ rather than a set of rules.

–

Joe Moran: A lot of the highway code is about sharing space, it's about how to behave in a way that protects you, and also protects other people. I mean, I should say that even though a lot of it is based on kind of unspoken or tacit assumptions, it has been a success. I think the other thing about roads is that they are quite miraculous examples of social cooperation in a way, I mean, because they are dangerous places and people can die on them. Every car journey, you make death defying decisions unthinkingly and you let the driver in on the motorway or you flash your lights to let them in and those are kind of split second decisions that could all could go wrong at any time. And most of the time, they don't. My view about roads is that they turn you into not perhaps a socialist but they turn you into a believer, they turn me into a believer in government, I suppose, because they are, they are spaces that need to be governed. And they are spaces where you can actually prove that intervention actually works. But I can also see that for somebody else, there might be spaces where you believe in free will and freedom, because for the reasons I talked about before that the social aspects of roads are either invisible, or you forget about them in your kind of little sort of climate-controlled box.

–

Jack Stilgoe: Roads, and the rules of the road, have always been politically controversial. In America, when cars were new, pedestrians objected to these new machines taking over their spaces and creating new dangers. The car was clearly a massive boon to a society that was developing rapidly, but the car lobby was clear that others would need to get out of their way.

Peter Norton is a historian at the University of Virginia

Peter Norton: The courts and the authorities and the experts in the newspapers. And in general, public opinion was on the side of the pedestrian, so much so that there was talk about restricting cars, making them incapable of going faster than 25 miles per hour and otherwise controlling them so that pedestrians rights would continue as they had. And the people who wanted to sell cars, though, saw that as a threat. And they organised to argue that it's a new era. It's a new age, it's the 20th century, it's the motor age, and therefore, things have to change. And pedestrians' rights are a throwback to another era. And today, it should be priority consideration for drivers and pedestrians were legally restricted. Really, by the 1930s it was is ubiquitous in the US?

–

Jack Stilgoe: If you believe the self-driving car enthusiasts, we are now on the cusp of another new age. Another new technology is making inroads into the way we move about.

Alex Kendall runs a start-up called Wayve that is testing self-driving cars here in London, which is very different from the not-so mean streets of Phoenix.

The fact that London is a hard place for a computer to drive is precisely why he thinks it's a good place to put his vehicles through their paces.

For Wayve, the challenge is to make learning to drive into Machine Learning

–

Alex Kendall: Our driving school is London. So we have a fleet of vehicles that we test every day, through the complexities of London and other cities within the UK, to learn how to drive. And we're really taking machine learning approach. So we let the vehicles learn to copy how humans drive. We've seen machine learning solve problems that are previously unthinkable for humanity to solve, like image recognition, natural language processing, or things like game playing agents, we've seen machine learning beat the world champion at go one of the hardest board games that humanity knows. So these kinds of these kinds of problems are really interesting and draw parallels with the problem of autonomous driving. And when you start to apply this kind of massive scale

data and learning, you start to see some really intelligent driving behavior emerge. And that's what we've been able to demonstrate on the roads in the UK.

–
Jack Stilgoe: He does however, admit that driving is not very like playing a game. Rather than knowing the rules in advance, machine learning is looking for patterns.

Alex Kendall: If you think about the game of Go, that's state space is fully observable, you can see every tile on the board, as you say the rules are predefined and well known. But in autonomous driving the state space is very noisy, it's not fully observable, you have occlusions, you have unknowns, you have sensor noise, all of this kind of stuff that makes the decision-making problem so much harder. I'd also maybe move us away from the concept of a rule but more patterns in the data. If you take a more chaotic driving scene, I don't know, like a place I visited, like Ho Chi Minh city in Vietnam. You could argue that there are less rules, not followed driving behaviours like we might see, in London, where a traffic light that's red usually means stop and things like that. It's a bit more chaotic. I would still argue that there are very, very strong patterns in how people behave.

Jack Stilgoe: A lot of self driving technology is being tested in very neat and tidy environments like in Phoenix, Arizona. Is London then a harder case? Does that make it a bigger challenge? Or does that make it an opportunity, or what?

Alex Kendall: It's an absolute opportunity. And I completely agree, this is what most people most commentators on the autonomous driving industry completely miss. It's a completely different game going from Phoenix to London. If you think about Phoenix in Arizona, it's a city that's pretty much always sunny. I think there's something like 30 rainy days on average a year. The roads are wide boulevards, that grid-like streets with very well structured lanes. We compare that to London. With its mediaeval heritage, there's about 100 rainy days a year, it's got about 15 times the population density. Similarly, increases in the road complexity with roundabouts, unstructured lanes, bus lanes everywhere, intersecting lanes, I mean, the kind of stuff that that you'll be familiar every day on the roads in London. If you look at those different factors around population density around with a complexity around roading complexity, it's about a 30 times harder environment to drive.

–

Jack Stilgoe: It’s about scale - equipping the software not just to drive in the easy places, but to drive anywhere. Wayve are sceptical of Waymo’s approach, which we heard about in the last episode, that relies on exquisitely detailed digital maps

–

Alex Kendall: So what I call the AV 1.0 the first generation approach really relies on building a high definition map that tells the car, where to drive through 3D space and where to look for things like traffic lights or road signs, so that it can really be told how to behave. If we think about last year, when the pandemic started, we had, we had all the footpaths within London widened to allow for social distancing. And this extended the footpaths through barriers everywhere. Now, overnight, that would invalidate your high-definition map and cause your system to go out of operation in the first generation of autonomy technology.

–
Jack Stilgoe: He talks about his cars learning to drive as though they were teenagers sat there with a driving instructor

–

Alex Kendall: From a product perspective, it has a very similar mode of operation to humans, it learns to drive from computer vision, and learns from experience and feedback as well as self supervised learning from observing the world like we do through our whole lives. And it drives from a basic Sat Nav map, like you might get from your satellite navigation system. So in that sense, it is based on a very similar scalar model, a scalable model of driving as humans. But of course, there are there are many pros and cons between a human that can drive and an artificial driving intelligence if we want to dig a bit deeper, but in essence, yes, we're learning to drive like a human.

We have a fantastic group of safety drivers. And they're really part of our development team. In fact, a number of them used to be human driving instructors. And what's neat about that is that our culture is really to look to find the interventions and find the things that the model doesn't know and teach them those skills. Look, we've learned to drive through roundabouts, traffic lights, multi-lane manoeuvres, and things like that. And we've learned a system that can really understand the complexities of these roads. I mean, just if you go out of our office and around the block, you're immediately thrown into a, a very complex roundabout, I think there was some roadworks put there the other

day that that made it even harder, you've got unprotected turns everywhere, you've got all of this rich, unstructured environment right on our doorstep.

–

Jack Stilgoe: But teaching a computer to drive is always going to be different from teaching a human to drive, using their common sense, through a world that has been designed for them. I asked him what his car would do at a zebra crossing

–

Alex Kendall: That's such a good example. Because I would argue that the road code as written is a necessary but not sufficient set of information to know how to drive within, let's say, the British driving culture. So much of it is through common sense and understanding of that driving culture. I think, if you were to try and write down those rules as a human, you would really struggle because do you set thresholds for if there's a human within x metres of the crossing? Do you stop? Do you try and predict where they're going to move? That's, that's inherently noisy. I mean, if you try and construct a rule set, as you get to more and more edge cases, and I could, I can't give you some weird and wonderful ones we've seen in this kind of scenario, your rules are going to become more and more contradictory. We let the machine learning model really figure out what the patterns are the underlying concepts and the data that led to that. And that's what's really powerful because at scale, we compound at scale, we don't become brittle, because we get more and more data that lets us tease out those underlying concepts through machine learning.

–

Jack Stilgoe: Letting the machine figure it out might be rather nervewracking if you’re a pedestrian who’s used to catching the eye of a driver before crossing. It’s a complicated problem. Maybe it would be much easier to change the roads to make things a bit more straightforward.

In the last episode, we heard from Lucy Suchman, professor of anthropology at Lancaster University. In 1987, she published a hugely influential book called Plans and Situated Actions, based on years of research at Xerox. This challenged the idea that machines could be made to think like humans. Just as humans exist in a wider world that they have built to suit them, so machines need to have systems around them that help them do their jobs.

–

Lucy Suchman: If we, if we accept the idea that in order to make these systems safe and effective, we have to, we have to close the worlds around them in necessary ways, then then we should listen to all of these proposals and all of these promises and imagine, okay, what kind of closure, what kind of enclosure, what kind of redesign, re-engineering of the environments in which they operate would be required? And what do we think about that? Do we think it's good? Do we think it's, it's practical and reasonable? So that's one question, what's the actual proposal for for not only the engineering of the autonomous vehicle, but the re-engineering of the world that it's going to be operating in? And then the second question is, what are the political, economic vested interests in this proposal for our futures? Right? Who benefits? Who loses? What are the costs? What are the lost opportunity costs? So if we're thinking about urban transportation, in what way does the focus on autonomous vehicles? who's interested in that? Who's promoting that? In what way does that focus marginalise other discussions we could be having about the future of transportation that would be oriented to cutting down on fossil fuels, supporting more pleasant and effective pedestrian, bicycle, public transport, electric buses, whatever, you know, whatever other kinds of future visions there would be, where the independent, you know, autonomous car is not a player. You know, people like Elon Musk, and others with vested interests are never going to encourage us to think along those lines. So we have to take everything they propose with a grain of salt, recognising what kinds of interests are driving that, and make sure that we're making space for radically different forms of innovation that we might be thinking about, than the ones that are being put forward. And, and it's hard, it's hard, because those people have very loud voices, you know, they they command enormous resources.

–

Jack Stilgoe: Alex Kendall accepts that there may need to be some give and take, but he doesn’t want to dwell on that.

–

Jack Stilgoe: So and then it might make sense to, in effect, say, you know, if we want to realise the benefits, then, you know, we might need to change the world in some way.

Alex Kendall: One more thing, though, is that the vehicle itself needs to have the embodied intelligence, the onboard understanding to be able to make its own decisions. And all of this kind of signal that can come into it from external connectivity, should be able to aid that. But if it sees if it's told to go straight, but it sees a cliff in front of it, the car should never be compelled to drive off that cliff.

–

Jack Stilgoe: If something does go wrong, like when the self-driving Uber killed Elaine Herzberg, and we're relying on software, would we know why a self-driving car did what it did? Machine learning people call this the problem of interpretability. As machine learning gets more and more complicated, interpretability may get harder and harder

–

Alex Kendall: Trust primarily comes from performance and safety, which is what we need to, to what we really need to drive first and foremost. So at the extreme, would you take 100% performance system that 0% interpretable? Or would you take 100%, interpretable. system that's 0% performance? I think everyone would choose the performance system. But having said that, I think that interpretability can really aid a number of things from human understanding to even engineering development of a system. It's worth comparing to things like, you know, the airline industry. Does every passenger understand why and how a plane works, what it's doing at each stage of that mobility service? Do passengers expect it? I don't think they do. But if you think about regulators, if you think about engineers, if you think about pilots, I think on this side, they do expect that. And so you can kind of draw some parallels there and I think there's going to be a time varying function to how society sees this, that maybe there might be a higher expectation to begin with. The other parallel, it's interesting to draw is with humans. When humans make mistakes, often they come up with counterfactual or or post ad hoc reasons for why a decision was made. Was it the actual causal reason that happened during the actual real time behaviour? I don't think that's always the case. And so what we see is, is humans need to be able to reason and give a plausible explanation for what they've done or at least give evidence that we might be able to resolve it. I think we can do better with machines, I think we can find maybe not causal understanding but certainly correlated. We can draw correlations as to why behaviours were done and we can certainly put in place things to improve the failure mode that we've seen.

–

Jack Stilgoe: The optimism is boundless

–

Jack Stilgoe: Do you have a sense of, you know, what the technology can't do at the moment and may never be able to do Is there a sort of limit where we can just say, actually, they'll never be able to drive in exactly this way? So we need to be honest about that.

Alex Kendall: Fundamentally, No, I haven't seen anything, any edge case, any scenario that I've been 'that's just intractable'. What we've seen again, and again, and again, over the last decade is given the right amount of quantity and clean and, and right type of data, given enough data, given enough compute and given a sufficiently large neural network, we've been able to resolve any problem.

–

Jack Stilgoe: But for all of the optimism, it's still not clear whether a car without a driver solves London's biggest transport problems. Bruce McVean is acting assistant director for city transportation at the City of London Corporation.

The City is the business district at the heart of London. It's the oldest part of an old city. Before the pandemic, more than half a million people commuted into the City every day, most of them walking, cycling or using public transport.

–

Bruce McVean: What's interesting from an autonomous cars point of view, is, you know, central London, like other central areas, and cities is not somewhere where we really want lots of people driving, or being driven, I guess, in the case of autonomous cars. So we're lucky that we have very high levels of public transport use. And we also have growing numbers of people cycling, and walking is by far the main way that people already travel around the area. So really, those are kind of the ways of travelling that we want to protect, and encourage more of. So I guess for autonomous cars, not so much for what it means, but one of the things that we know we don't want is we don't want it to lead to an increase in private or private use of shared vehicles.

–

Jack Stilgoe: For a dense city like London, it often comes down to a question of space.

–

Bruce McVean: I suppose does it make a difference, if all we do is switch those out to autonomous vehicles and get, you know, some of the advantages that would come from that, like I said, you know, safety is a big one for me, I think of one of the kind of really interesting opportunities here is around how we how we eat or how autonomous vehicles in in due course, and, you know, once the technology is proven, can help make our streets, hopefully make our streets safer places. I guess the thing for me is, we know that, that the biggest challenge that we face, I guess as transport planners, and it's the same for most urban areas, and particularly central urban areas, is we're dealing with a finite amount of space, you know, we're not creating more streets. It's not unusual, I guess, in Europe is, you know, we are as a separate kind of where London began. So we're a historic area, and our street network, although the buildings that line it are increasing tall and modern and everything else, you know, our street network is largely the network that we had in mediaeval times.

–

Jack Stilgoe: In the City of London, they know what they’re trying to do, so they ask whether technology is going to help them solve their problems.

–

Bruce McVean: I think in some way, like London, you know, you're you've got a, you work with the assets you've got, and we've got really good assets of a really good public transport network, a growing network of safe cycle routes. And, you know, a very walkable city, particularly in the centre. If self driving vehicles appear, then they've got to support those assets and not compete with them.

–

Steve Gooding runs the RAC foundation, once part of the Royal Automobile Club. He was formerly a policymaker at the British Government’s department of transport

–

Steve Gooding: I'm actually very interested in why driverless is such a big thing in the automotive sector. What is it about this technology that the auto companies think drivers want, because they clearly must think it's something that we want, otherwise, they

wouldn't be pouring billions of pounds into it. Billions of dollars, billions of euros. And clearly they are.

The road space in the centre of London is contested space. There are pedestrians, but not just pedestrians, there are office workers coming out of their offices at lunchtime, they've got their earbuds in, they're thinking about their job, they're thinking about going and getting a sandwich, you've got tourists all wandering around, looking up a whole bunch of people not looking where they're going. A whole bunch of other people in a hurry. So we have people making deliveries, we have cycle deliveries, we have motorcycle deliveries, we've got vans whizzing around, it's a very difficult environment, to see how a vehicle is going to navigate that safely, without actually being extremely slow.

I always use the example that in her later years, my mother was quite frail. And if the mini cab driver hadn't been able to get out of the cab to help her in, she wouldn't have been able to get into the mini cab. So it's one thing to have a vehicle that can drive itself, but if the passenger can't actually get out of the house, and into the vehicle, they aren't going anywhere. That's why we need to think about this, not just as a vehicle doing something, but as a whole system providing a service.

Some time ago, I said at a conference that I sometimes felt that the government's enthusiasm for making this country a testbed for driverless technology would give us all the benefits that Christmas Island got from being a testbed for nuclear weapons. You get all the pain, but you don't necessarily get much of the gain. Now what I meant by that was there is value to be had in making this country a testbed for cutting edge automotive technology. Without question. We attract the best brains, we attract inward investment, that's great, but it's inward investment into the development of the technology. The trick is both in making sure that technology is adopted and adopted in a way that's of benefit to the GB economy and to our citizens, and that it's adopted and put into production by companies that are working here and employing people here and the question mark I think I and others have still got is, are we going to be able to make that jump?

Are we going to be able to say, well, because these vehicles were tested out here, because this technology ran around Oxford, and because some people are doing trials in London, that's good? Or is there a risk that the companies that are doing that are proving their vehicles and their technology in a very difficult environment, but with one eye to selling that, to setting up the manufacturing, to delivering the benefits elsewhere

in the world, and I think that's something that the Centre for connected and autonomous vehicles, with the Department for business hats on, still need to work out?

Jack Stilgoe: There is still a big gap, it seems, between what the technology can do and what the world needs. Peter Norton again.

Peter Norton: One of the most constant frustrations for me, is to read reports, often by the people interested in promoting driverless cars, that they'll make everything better, including for pedestrians and cyclists, because of the fact that an algorithm will be controlling the car and it will, it will not take chances, and it will not, you know, threaten pedestrians. But I think that's a naive view. And I think that what it lacks what makes it a naive view is the neglect of experience. We know what happens when we have a conflict between a new profitable vehicle like a conventional automobile was 100 years ago, or like an autonomous vehicle is today, namely, the people who stand to gain from selling these things, will ensure that they get the priority consideration that the success of their business requires. Now 100 years ago, what that meant was changing laws to sort of compel pedestrians to cross the crosswalks and otherwise concede their right to the street to the automobile. Well, we could see analogues of that for autonomous vehicles too, because, you know, if you're in the business of providing autonomous vehicle rides, you won't make money if your vehicle is stopping every three or four seconds for another pedestrian. So what you'll do instead is you'll organise like they organised 100 years ago, to get the kind of change necessary to keep the priority on the vehicle.

Jack Stilgoe: It seems, ultimately, that we can never ignore the politics of roads - Who gets priority, who has the right of way?

In this podcast, I've been exploring what self-driving cars might mean for society. The question is often asked 'when will self-driving cars arrive?' I think that's the wrong question. The real question is WHERE will self-driving cars arrive, and in what form. We could see self-driving cars stuck in the few places that suit them. Other places could say no to the technology or they could adjust the world to realise its purported benefits.

Maybe Zebra crossings are just too unpredictable? Maybe we need traffic lights instead? Maybe everyone needs to stay in their lane, with pedestrians keeping to the pavements so that humans and robots can move more efficiently and safely through our cities. Or maybe we will all just get used to sharing our roads with robots, slowly learning what they can and can't do.