

2017 Meeting of the Minds - 'Driverless Urban Mobility: The Path Towards an Autonomous Future' - Dr. James Kuffner

April 14, 2017

[2017 Meeting of the Minds Webinar](#) – “Driverless Urban Mobility: The Path Towards an Autonomous Future”
March 28, 2017

Dr. James Kuffner, Chief Technology Officer, Toyota Research Institute

Transcript:

It's my great pleasure to be speaking to everybody today.

I'd like to talk a little bit about some of the projects and technology that we're working on here at TRI, how the industry is moving towards this technology and reflect a bit on how it could potentially transform our cities in profound ways.

I got my start working in robotics and worked on a lot of different humanoid robot platforms, both in Japan and in the U.S., at the University of Tokyo and then Carnegie Mellon University.

In 2009, I went to Google and started working on their self-driving car team. I worked on a lot of AI and robotics related projects there.

Early in January 2016, I joined the Toyota Research Institute under our CEO Dr. Gill Pratt, who many of you may know from the DARPA Robotics Challenge. TRI was established to partner not only with leading research universities – including MIT in Cambridge, Stanford right here in Silicon Valley, and University of Michigan in Ann Arbor – but also to be a place that partners with companies that are thinking about the future of this new technology as it transforms both vehicles and robotics.

Our key focus areas are in automotive safety and autonomy.

We're also looking at mobility, especially where it relates to an aging society.

We're also looking at how cloud computing machine learning could accelerate scientific discovery. For example, could you explore new catalysts for batteries that would lead to a much higher energy density or find a new material that is twice the strength of high-tension steel but as light as carbon fiber? We're looking at all of those technologies.

Today, I'd like to talk a little bit about innovation as represented in history, and how it can be used as a model for how we can expect the maturing of this autonomous vehicle technology. I'll talk about how it could impact future cities, and reflect a bit on how that relates to my favorite topic, which is cloud robotics, and how it all ties together.

First, I'd like reflect upon the computer.

Many of us consider the computer as the most complex machine that humans have invented. In fact, today's computers cannot be designed without a computer.

If you assemble many of them together, you get the modern data center, which has truly transformed the way that we do business and store and manage data, especially now with all of our connected devices.

I would argue that the robot is really the most complex machine that we've invented because you not only have a computer, but also actuation and motors, and you have to obey the laws of physics and real time constraints in all of the computation that you're doing.

If computers manage bits of information, robots are managing and organizing atoms and applying the organization of information to the physical world. So what happens if you have the most powerful computing resource, which is the modern data center, connected to a robot? That's happening, and it's very exciting, both in the area of cars as well as robot machines.

If we think about the evolution of the automobile, there were gas powered engines back in 1885. Over the next 30 years, there were incredible advances in the transmission, engine design and manufacturing at scale that led to the first widely affordable automobile, which completely transformed the society.

The same thing happened with the computer. There were early vacuum tube-based prototypes in the 40s and 50s. They were very expensive and only a few research institutes and government labs around the world had them. But, there was incredible advancement in the hardware, the storage displays and the development of transistors and integrated circuits that led to a widely affordable personal computer that also completely transformed our society.

The shift that all of us have lived through is the smartphone.

Back in 1983, you could buy a Motorola DynaTAC for about \$4,000. It sometimes completed a call.

Over the next 30 years, incredible advances in the network and transmission speed, cost, size, weight and power led to the smartphone, which now has overtaken desktops worldwide as the main computing platform for people.

What's truly remarkable is that we all take for granted the fact we're carrying a supercomputer around in our pocket that's more powerful than the most powerful supercomputer of 20 years ago.

When we think about intelligent robots and cars, I like to count the clock from when Honda released their first fully self-contained walking humanoid. That means it had onboard power and onboard computation.

Over the last 20 years, we've had incredible advancements, not only in the sensors with depth cameras and LIDARS, but harmonic drive gears, better actuation, control planning and machine learning that are really starting to make robots incredibly capable.

It's always dangerous to speculate on when they'll become ubiquitous and affordable, but I certainly hope, based on the historical precedent, that it will happen in our lifetime.

I'll talk a little bit about autonomous cars because there's been incredible advancement and excitement about this technology starting to mature and really starting to hit the industry in a big way.

Early advancements in autonomous driving were primarily led by military investments. DARPA had been giving

away and funding research at universities for decades until the DARPA Grand Challenge in 2004. The Grand Challenge flipped the problem around. Instead of just funding individual projects, DARPA made a prize bounty and let people crowdsource the ideas and assemble projects to try to win a challenge.

The first challenge in 2004 was essentially GPS Waypoint following. There was no winner, but the Carnegie Mellon Sandstorm had traveled the farthest. The next year they held it, and there were actually a couple of teams that finished. The top prize went to Stanford Racing.

Two years later they did the Urban Challenge. They upped the ante with a \$2M prize, and there were several teams that finished the complete route. The top prize went to Carnegie Mellon that year.

That led to a project being formed at Google made up of a lot of people who have been involved in the DARPA challenges. I joined that team in 2009 when it started, and the goal was to try and explore whether or not we could make production quality code and cloud computing infrastructure in the development of something reliable that works well.

Today, there's an explosion of R&D activity. This headline from last summer reads “33 Corporations Working on Autonomous Vehicles.” There's now many more startups and companies looking at this technology.

There are lots of reasons why, and we can get into the reasons people are excited about this, but I think it actually comes down to how this technology can transform our society.

If I look at the ingredients for disruption, I think there's a confluence of strong partnerships that are necessary—partnerships between the universities (academia), industry and government. There is also a critical mass of talent and investment capital to try to bring an idea into the mainstream and bring products to market.

Let's reflect for a bit about this intelligent vehicle technology and think about how it may disrupt the design of cities. We take for granted that our cities fundamentally have been designed around cars. They've grown up around cars. The existence of freeways, traffic light systems, pedestrian overpasses and all of those things evolved because the automobile is a primary way for people to get around our urban centers.

What happens if you have autonomous cars and mobility as a service as your main means of transportation?

If we think about what it would take to enable a true driverless city, we look at transportation on-demand. Mobility as a service means that you may not have to worry about driving your personally-owned vehicle to where you work in a downtown urban center, looking for a parking space, and then having your vehicle sit idle and not be utilized while you're in the office all day.

Instead, the sharing economy will enable a much better utilization of that vehicle resource, and you can get where you need to go on-demand. Potentially, that could lead to a dramatic reduction in traffic, noise and pollution in our urban centers. If you just think about the capacity of lanes, there's a lot of urban centers that have curbside parking.

If every car could essentially drop you off where you needed to go and then go park somewhere else, you would, in some cases, immediately double the lane capacity and flow on many of the streets in an urban center.

That also means that land dedicated to parking lots, especially in urban centers, could then be converted to other uses—whether it's residential or commercial uses.

When we think about parking, it's useful to look at the current state of where things are at. It turns out that the average car is parked approximately 95% of the time in the U.S. with about 5% “on the road” time. Worldwide, people that live in urban centers spend an average of twenty minutes per trip just looking for parking.

The United States has about a billion parking spots, but there's only about a quarter billion cars and trucks, so that means there's about four times more parking spaces than vehicles.

A study that was done for Los Angeles in 2015 found that there were about 200 square miles of land currently dedicated to parking in the county. That is about 18.6 million spaces. From the data available, that's approximately 14% of all the land area in Los Angeles County.

What would it mean if you can be a lot smarter with not only about the utilization of the seats in your car, but also utilization of the parking spaces? When we think about parking garages, what would driverless cars enable in terms of the space utilization of a parking garage?

First, it doesn't need to be in the urban center. It means that a lot of cars could essentially be like a valet, dropping passengers off where they need to be in a downtown center and then parking itself outside of the center. Essentially, valuable land that's currently dedicated to parking structures in the urban centers could be rededicated to other uses – whether it's commercial or residential.

It also means that the parking structures can be designed to be more efficient and densely packed. Let's say there's a robotically-managed parking garage. It may not have typical requirements for parking structures that contain stairs or elevators or wide alleyways for pedestrian access.

The data driven dispatch of on-demand transportation could also enable the system to do dynamic load balancing, matching the vehicle supply in the areas where the vehicles are needed according to demand patterns. This is a great opportunity to do data mining and machine learning, looking at the transportation demand curves for all of the people, and then essentially running a giant optimization problem that will match the location and utilization of the fleet with the transportation demands of the city.

It also means that parking lots can do a lot more than just be a place to store the car. You could centralize charging stations and replace traditional gas stations. Not only would you be able to centralize refueling and recharging from the currently very widely distributed land in urban centers, you can also reclaim and recover the land that's currently dedicated to gas stations. Centralizing refueling into a single location can also be environmentally clean with solar panels on the roof generating power to charge clean energy vehicles. And, in the case of fuel cells, these central refueling depots could have fuel cell vehicle refueling located in these structures.

Additionally, while a vehicle is being stored, it can be maintained. In the case of self-driving cars, there's a lot of sensing equipment that needs to maintain calibration. You could imagine that calibration, health checks and routine maintenance for the fleets of vehicles could be done in one location, making it much more efficient.

People have talked a lot about the number of vehicles out on the road and how mobility as a service could affect traffic patterns and peak demand. Right now, most of the vehicles on the road in the U.S. – unfortunately during commute hours – over ninety percent have a single passenger riding in them – it's the driver. There's also a lot of delivery trucks with just a single passenger in them.

Imagine if your personally-owned vehicle could park itself in a recharging station that was centralized, and any shipments from online orders would automatically get loaded into your trunk of your car while you were at

work. And then when your car was dispatched to pick you up, you already had all your packages that you had ordered delivered in your car. That could actually remove a lot of the delivery trucks that are currently on the roads during daytime hours.

I think this could open up new ideas to have logistics hubs and parking as a new way to think about the on-demand economy.

Bigger picture...since current cities have been designed around manually driven cars, we can think over the longer term about how urban design could be changed with automated cars and autonomous technology.

I mentioned converting parking lots and gas stations to green spaces. You could also potentially convert curbside parking to bike lanes or wider pedestrian access. I think this could make our cities not only more livable, but even potentially safer. Having a lot of parked cars on the sides can often lead to blind corners, which is the source of many accidents.

We can explore a truly radical urban center redesign. I'm a big fan of the "FuBgangerzone" – the pedestrian zones that are in a lot of European center cities. The reason why they're so wonderful is because there's a lot less noise, better air quality, and they're simply a lot safer to just walk around cobblestone streets in the downtown center. Sometimes, it's done in Europe for a practical reason because these cities pre-dated automobiles, built when horses were around. Certainly, it is much nicer to enjoy a city center without the worry about a lot of vehicle traffic.

Imagine if you simply moved all vehicle traffic underground. In some cities in Russia, they're moving a lot of shopping underground. Part of it is for weather reasons because half a year it's extremely cold. But, they've been able to essentially pre-pour large concrete walls and then excavate the spaces in between them to build giant underground shopping malls in Russia.

When we think about this technology, we can also think about how electric vehicles and zero emissions vehicles will have a huge positive impact on our cities.

For new mobility, Toyota is experimenting with a lot of different new form factors. One of them is the i-Road. It's an electric vehicle design that has active balancing and stability.

What's really wonderful about it is that its width is such that you can fit two of them side-by-side. This essentially doubles the number of lanes that that you have with the existing streets without changing or doing construction to add more lanes. Imagine doubling the throughput if you had urban centers with these types of vehicles.

These can seat up to two people. It's nicer than a motorcycle because of the closed cabin with climate control and protection from wind and rain and weather. Also, in one parking space you can fit up to three of these. So, you could imagine not only saving parking and relocating it from urban centers but you can have much more densely packed parking because these vehicles are so much smaller.

These are still experimental, but it's interesting to think about what form factors will work well, especially if we do have a lot of mobility as a service/transportation on demand, and a lot of people as single passenger riders.

For fuel cell vehicles like the Toyota Mirai, Toyota is making its fuel cell patents available. In California, I had a chance to test drive it, and it's really wonderful with a long range and truly zero emissions. Wheel to well, it's very environmentally friendly and refuels in just about five minutes. What's wonderful is that these are perfectly

well-suited when you have central parking structures or refueling stations. Fuel cells can also be used for large trucks and buses that are currently diesel-powered.

I want to talk a little bit about how all of this technology fits with the parallel advancement in artificial intelligence and machine learning. This part of the reason why people have been really excited.

If you think about search-based planning, back in 1997, Garry Kasparov, who is the reigning World Chess Champion, was defeated by IBM's Deep Blue. IBM's Deep Blue system searches about 200 million chess boards per second. That's not how a human plays chess, but it is amazingly effective.

Moore's Law, which has been the doubling of CPU power every 18 months, has been remarkably consistent and has been allowing more and more brute force types of algorithms to become more and more effective and solve more and more problems.

About a year ago, the team from Deep Mind, as part of Google, created AlphaGo, which defeated the world Go champion in Korea. Many AI researchers didn't believe that this was going to happen for many, many decades.

So what was the big change? It's machine learning.

I think even more, it's the ability to do massively parallel computations in a cloud in modern data centers.

We think about what will happen when we have these data centers connected to a lot of devices, not only our smartphones but our robots.

We had started a project when I was at Google called the Cloud Robotics Team, where we talked about this concept of turning your old smartphone into a connected robot. A lot of people were turning over their smartphones every 18 months and essentially giving them to their kids to play games. These are essentially supercomputers, and they have all the things necessary to make a robot.

We open-sourced a lot of that work and ported the robot operating system from Willow Garage to run on Android so you could turn your old phones into a robot.

The larger picture is that not only is Moore's Law powering incredible advances, but wireless networking has enabled a remarkable transformation of computation and data on demand.

If you think about the peak mobile broadband speed that we've seen in the last 10 years, there's been a 1600x growth. Nobody would have believed 10 or 15 years ago that people would be able to watch high-definition video on their mobile devices, and everyone could watch it all at the same time. Incredible bandwidth has enabled a lot of new applications.

When we think about cloud enabled robots and cloud connected cars, I think one of the key advantages is that it provides a shared knowledge base and unifies information about the world in much the same way that humans use Wikipedia or other websites that organize information as reference materials.

You could imagine cloud-connected robots and cars going to central databases to find information about things in the world. It also means you could offload a lot of these heavy compute tasks to the cloud. Instead of having a car or robot that needs to carry around all the compute power and storage necessary to solve a problem, it'll be able to use this computation and storage on demand as a service.

That means, there can be fewer needs for hardware upgrades that can happen invisibly and hassle-free in the cloud, as well as fewer software pushes. As soon as a software push is patched in the cloud service, all of the robots are connected. Cars have an invisible hassle-free patch.

I think this idea of a reusable library of skills, when we think about motor skills for an intelligent robot or driving behaviors for cars, we can think about data mining the history of all of these cloud connected machines in order to make all of them more proficient.

Just a couple of quick examples. One has already been transforming the way that we interact with our smartphones. We are now able to have real-time speech to speech. That means you can speak in one language and have your mobile device translate in real time to another language. Then have someone else speak back in that other language and then have it translated back to you. It's the Star Trek communicator that we've all been dreaming about.

There's no reason why every cloud-connected robot shouldn't be able to speak all languages.

Another example is perception. Many of you may have used, years ago, the experimental app Google Goggles, which allows you to point your smartphone at a famous painting or a statue, and then it would do a visual search and be able tell you who painted that and when it was painted.

There's no reason why we couldn't have a robot that is able to capture an image of something in the world and then retrieve semantic information about it – how to grasp it, how much does it weigh, what are the usage patterns and any other domain knowledge necessary for it to successfully manipulate that robot.

I'm showing you the human support robot which is a platform that Toyota is developing. It has been chosen as a standard platform for the RoboCup @Home in Nagoya this summer to try and challenge people to build reliable manipulation and solve some of these tasks.

The larger picture is it enables robot sourcing. Much like human crowdsourcing can help scale hard semantic problems globally like Wikipedia or map-making, the large-scale deployment of data-sharing robots and cars offer similar advantage.

I remember a colleague at Carnegie Mellon telling me "I know that this machine learning algorithm is going to work, it's just that my robot breaks down before I have enough data." I think, instead of having one robot run for 10,000 hours, let's try to have a hundred robots run for a hundred hours and collect the same amount of data.

I think that's the potential for the cloud robotics data sourcing.

When we think about connected cars and how this is transforming vehicle intelligence through deep learning, you think about LIDAR traces of cars and vehicles. It's already had a huge impact on natural language understanding, speech and object recognition with things like ImageNet.

But, the future of connected cars means that you're able to gather novel data and upload these new exemplars, and then update those trained models in the cloud and broadcast an updated model that's more reliable back to the entire fleet.

I think this is how we're going to scale some of the very, very hard problems in terms of perception and very, very tricky corner cases that are really difficult to deal with at scale.

Let's let the entire fleet learn and improve over time, and I think this is an exciting way that we could potentially crowdsource and robot-source the gathering of this data and improvement of the behaviors over time.

I want to end with one example of this which I think illustrates the point quite well.

There's a lot of companies that are investing significant resources in mapping. One of the unfortunate realities of a large mapping effort is that as soon as you gather the data, it's out of date.

How do you deal with the staleness problem of mapping?

This was demonstrated as a proof of concept back at CES in 2016. The idea is that many cars now have backup cameras or forward cameras that are able to look at images of the road surface. That means, they can detect where the line markings are.

Let's say you have a car that has a pre-image of the road surface and road signs and speed limits, and then is able to do a smart differencing. If lane markings get repainted, or if a pothole appears, those differential pixels are then sent up to a centralized cloud server which can then broadcast back to the fleet and stitch together an image of the roadway surface that can be kept up to date.

There's 10 million cars that Toyota sells globally each year. That means there's about roughly a hundred million Toyota vehicles running around the planet. Odds are that a Toyota vehicle will pass by a street near you every 30 seconds or few minutes. That way you could actually keep a map up to date without staleness.

This is really exciting and something that we're looking at and exploring. Going forward I think we're going to see an incredible diversification of these new application areas, not only for transportation, but in manufacturing, defense, medical, outer space and logistics' moving packages and objects around, as well as other consumer products as the technology matures.

In closing, I think it's very exciting to think about how the next generation of these intelligent vehicles will enable us to think and rethink how our cities and urban centers are designed. The promise of cloud computing and big data and the ubiquitous connectivity that comes with wireless technology is going to cause an incredible advancement in, not only the intelligence, but the safety, access and reliability of services that will be provided through these connected robots and cars.

Finally, when we think about the true promise of cloud robotics enabling these cheaper, lighter, smarter robots with the shared knowledge base, applying the robot-sourcing of data, data mining and machine learning to make the entire fleet more performant and more useful holds incredible promise.

I think one of the keys is to work closely with academia, government and industry and form partnerships to make things work. This is much the same way that partnerships stimulated the advances in autonomous driving dating back to the DARPA Grand Challenge which resulted in today's incredible explosion and exciting progress.

Those are the main thoughts that I wanted to share.

Thanks, everybody.