The autonomous car:

How Machine Learning broke through limits that Engineering couldn't get past

By Peter Keen and Ronald Williams

The autonomous car began as an opportunity that required breaking all kinds of limits: engineering, navigation, adjusting to traffic conditions, distinguishing objects, predicting what those objects might do, reacting in time, calculating quickly and juggling a vast number of ever-changing variables. The developers used more and more computer power to address these needs. But the initial bounding limit turned out to be very fundamental; *rule-based computers don't have pattern power.*

From the beginning systems with pre-programmed data definitions and structured procedures – rules – dominated the use of computers for everything from standard business data processing applications to even AI. And so it was logical that the early approach for autonomous cars was to define rules, derived from experts, and embed those rules in the software in the form "If..Then… Else…" "If the car is stationary and a pedestrian crosses ahead, remain stationary. Otherwise, check the brake is off ..If it isn't, then…" But that did not work, as the first Grand Challenge competition for autonomous cars made very clear.

DARPA (Defense Advanced Research Projects Agency) set up its Challenge, with driverless cars competing to complete a 150 mile course across the Mojave Desert for a prize of $1 million. None of the 15 finalist got more than 7.32 miles. Yes, the course was hard. According to the designer Sal Fish "It had rocks on it, left turns, right turns, dips, gullies, cactus all around the place. A drop-off on one side and then 5 miles later a drop-off on the other side. Barbed wire fences, animals that could come out of nowhere, train tracks. Anything and everything that a vehicle would encounter going through the desert."

But the terrain was not the real problem. Sebastian Thrun, who went on to head Google's self-driving project, observed that "These vehicles didn't fail because they weren't rugged enough. They failed because they didn't take in enough environmental information. None of them *saw* anything." There was no lack of accessories to see with; the cars were equipped with a combination of Lidar (laser rays), cameras, and GPS sensors. But the cars' rule-based computer systems could not recognize what they were sensing. There was no way to filter out noise and every possible situation had to be anticipated in advance and coded into the system. Thrun added that "Very few self-driving car people knew anything about machine learning at the time."

So 18 months later DARPA ran its second Grand Challenge, doubling the prize to $2 million. Once again, two hours before the race DARPA provided each team with a computer file of GPS coordinates for a series of waypoints, 237 feet or less apart, depending on the course terrain. This time, all but one of the 23 finalists made the 7.32 mark and five completed the 132 mile course. While there was still plenty of new and faster devices loaded on the car, the leaders emphasized that what changed was being able to really *understand* the wealth of data that the sensors were feeding back in order to navigate the course. Machine learning had entered the picture, breaking through the technology limit and in doing so shifting the entire direction of achieving the Ambition of a complete transformation of motoring from the physical car to the mental driver substitute.

In 2007, the Challenge was moved to an urban setting, a closed Air Force base in California. The 11 finalists out of the original 89 applicants had 6 hours to maneuver through 60 miles of the usual dynamic obstacles: intersections, and four-way stops, for instance. There were no pedestrians, bicycles or working traffic lights, but 30 Ford Tauruses equipped with crash cages and race seats driven by professionals added a little excitement.

Six of the eleven robot cars finished the course. The winning entry from Carnegie-Mellon University (Stanford was runner-up) was software-heavy, with four rule-based systems handling individual functions. The fifth one was the critical component: a machine learning model that builds up its interpretations of the multiple stream of input data from the car's sensors and fuses them to form a composite picture – the equivalent of a mental map – of the exact status of the car in its environment.

Without this machine learning capability, tests showed that the car could not handle even the fixed and static context. The rule-base systems could not cope with all the noise from the sensors - what is real and what not? As all the Challenge events showed, it was impossible to define and manage enough rules.

Google's leaders realized another limit. Given the speed of the on-board computer and algorithms required there was no way that they could succeed with an autonomous car if the car had to sense everything in real time. Maps were the key. If they could load the on-board computer with detailed maps of all the fixed surrounding, the car's system would only need to deal with what was dynamic. Two years before it officially began its self-driving car project, Google formed a team that included Sebastian Thrun, the developer of Stanford's winning entry in the 2005 Challenge and also noted for his work on robot mapping and computer algorithms for probabilistic reasoning, i.e., the foundations of machine learning as a computer science discipline.

It launched Google Street View, leveraging the sensor technology used in the Urban Challenge cars to record 360 degree views, along with high quality GPS data, as they drove down the street.

This pushed the limits of mapping and in doing so opened up more opportunity space for self-driving cars. Google's developers set the ambition of a "singular vision to photograph every street in the world." Thrun said that "Data was the secret sauce for getting self-driving cars to progress as well as they have. But it wasn't a matter of finding a data set and applying it. It was about creating the data set for that specific purpose. Street View wasn't a useful data set that was applied to self-driving cars. It was the output of the mapping exercise that made self-driving cars work so well."

Many industry veterans believed at the time that communication between the cars would be essential for self-driving cars. But these newcomers quickly came to understand that if the car knew where it was all times and what was happening around it, it would not need communications between cars

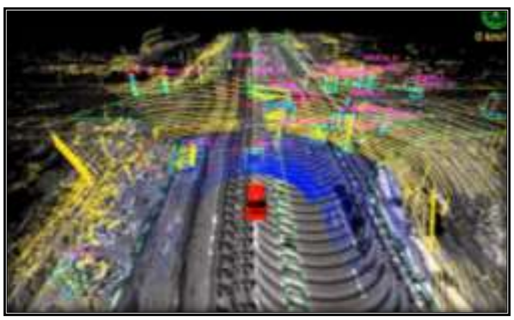

And so the Thrun team reverted to the centuries old art of map making, but this time at a level of detail not seen before. This image shows its use of data no human can ever see: the invisible light in the electro-magnetic spectrum, here pulsed laser rays processed through Lidar (Light detection and ranging). This is a superb map for a driverless car to navigate by and utterly worthless for a human driver.

Limits remained, of course – many, many limits with many, many dimensions and hidden gaps and traps. For the 2007 Urban Challenge, the only moving objects to detect were cars. But for the self-driving technology to be taken into the real world the car must be able to perceive not only a wide variety of moving objects, like the cyclist, policeman, other cars in the picture beside, but predict their behavior.

It would take a few more years for the more sophisticated machine learning algorithm of deep learning neural networks to provide the answers. Those years packed more and more data into the machine learning inputs. It took Google's Waymo 6 years to log 1 million miles. It is now logging 1 million miles every month. In addition, it processes 8 million miles of simulation a day. Today machine learning play a role in virtually every aspect of the self-driving car's software.

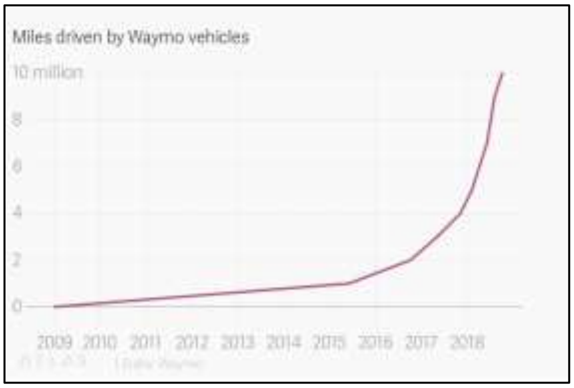

The impact of the learning data and accumulation of experience is captured in a comparison with GM. Waymo drove 350,000 miles in California in 2017 and had 3 crashes. For equivalent miles, GM, the consensually rated second in the development competition, had 58. Yes, that is a 20:1 ratio and it illustrates one of the most critical factors in AI: the advantage goes to the player with most data and most training trials.

Time and timing are always factors in technology innovation. The issue for every executive looking to set the firm's Ambition is "Is time on our side?" Is it better to wait and let Waymo's AI experience be a guide to our actions, when we are ready (Morph rather than look be Break Away and be a pace-setter)? Do we lack the data now but can get it when we need it?

Or are we self-limiting our opportunity by delaying the transformation opportunity? Even for the leaders, there is a long way to go in breaking open the practical opportunity limit. Waymo defines this as "90% done and 90% to go."

The 90% to go refers to all the unusual situations yet to be learned. The machine learning systems will continue to adapt and improve with every mile driven and every mile simulated. And the machine learning algorithms will continue to be refined and extended.

Humans have an amazing ability to perceive and understand what is happening around them and handle it. But what they understand and how they react is based on what they have experienced. Any one human has a limited set of relevant experiences. Humans get distracted, too often texting while driving – machines don't. Some cannot see that well or have reduced reaction time. As in all of AI, there is a complex interplay between human and machine limits; it's a matter of cognition – perceiving, understanding, applying – and processing – brainpower, speed, memory, reaction time.

The self-driving car was unable to match even basic human capabilities for the first DARPA Challenges. Engineering and rule-based software didn't break through that limit. Machine learning is doing so and the evidence is that it will create new levels of performance and go beyond humans' limits on what they see, how much attentions they are paying and the situations they have experienced to eliminate accidents.

Then, the opportunities will be, if not unlimited, expansive and far-reaching in potential. Autonomous vehicles will likely break the limits on:

Someone having to drive
The need to own a car
Mobility for seniors who can no longer drive
The last mile of package and food delivery

The need for parking garages
The risks and costs of insuring drivers
DUI and DWI frequency and impacts
Where we live
And many more unforeseen impacts on everyday life.

90% done or 90% to go?