Road Testing or Simulation? – The Billion-Mile Question for Autonomous Driving Development

By **Dr. Luca Castignani, Chief Autonomous Strategist, MSC Software**

Figure 1. Autonomous Driving is a Fast Growing Industry

Just two years ago, the future looked so bright for autonomous vehicles and everything was coming at a fast pace. Most auto OEMs, Tier 1 suppliers and hundreds of start-ups around the world presented their aggressive plans to bring self-driving vehicles on our roads. Traditional OEMs were a bit more cautious, but new players were very bold in their announcements, which is understandable considering that they had to convince their investors that the future of Mobility-as-a-Service (MaaS) was coming very soon.

And then, something tragic happened early in 2018. Something that had not been foreseen by many people, blinded by the hype that was burning millions of dollars on a daily base. A vehicle used by Uber to perform self-driving test (it would be wrong to call it "a self-driving vehicle", since nothing like this exist as of today, and what we see on our streets are just "test platforms") hit and killed a woman crossing the street in Arizona.

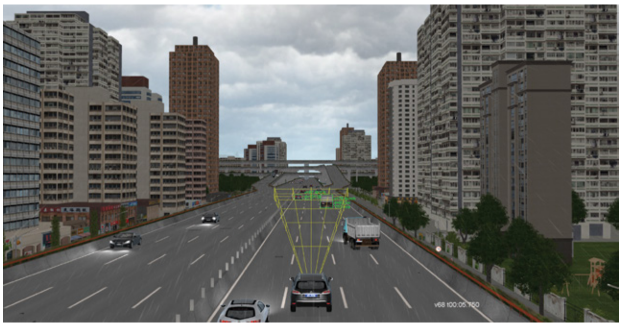

Figure 2. Simulating Autonomous Vehicle driving on a highway in Beijing, China. Simulation done in VIRES VTD.

The reaction from the public was immediate and strong, which put into question not just Uber but also the whole self-driving effort (NVIDIA lost 10% of its market value in the following 2 weeks). The Washington Post titled "Fatal Uber crash spurs debate about regulation of driverless vehicles" [1] and the Guardian "Uber crash shows catastrophic failure of self-driving technology" [2].

What should be the lesson learned from this case? A LinkedIn user expressed it in the best way: "It is completely unacceptable that undesirable beta software is being tested on the roads. This is not an online game where you have several lives".

Achieving autonomous vehicle functionality and safety requires millions of tests to cover all driving scenarios, and there is no way to get even close to that if not extensively using (to an extent never dared in the history of engineering) simulation in the virtual world. Trying to achieve level 5 autonomy only (or mostly) with road testing is as useless as boiling the ocean.

We are engineers, so let's do some basic mathematics. According to the US traffic accident reports [3], statistically, there is one person killed on the road for every 140 million kilometers driven. Therefore, to statistically prove (with 95% confidence) that an Autonomous Vehicle is as good as a human driver, it has to be tested for 415 million kilometers without causing any death [4].

Many self-driving enthusiasts claim that "autonomous vehicles will reduce the number of people killed on the road by a factor of 20" [5]. But this a very bold statement that needs to be proven before being accepted by the general public, regardless of how appealing it sounds. After all, we are scientists, and we believe in data. And the data that we need to prove that the

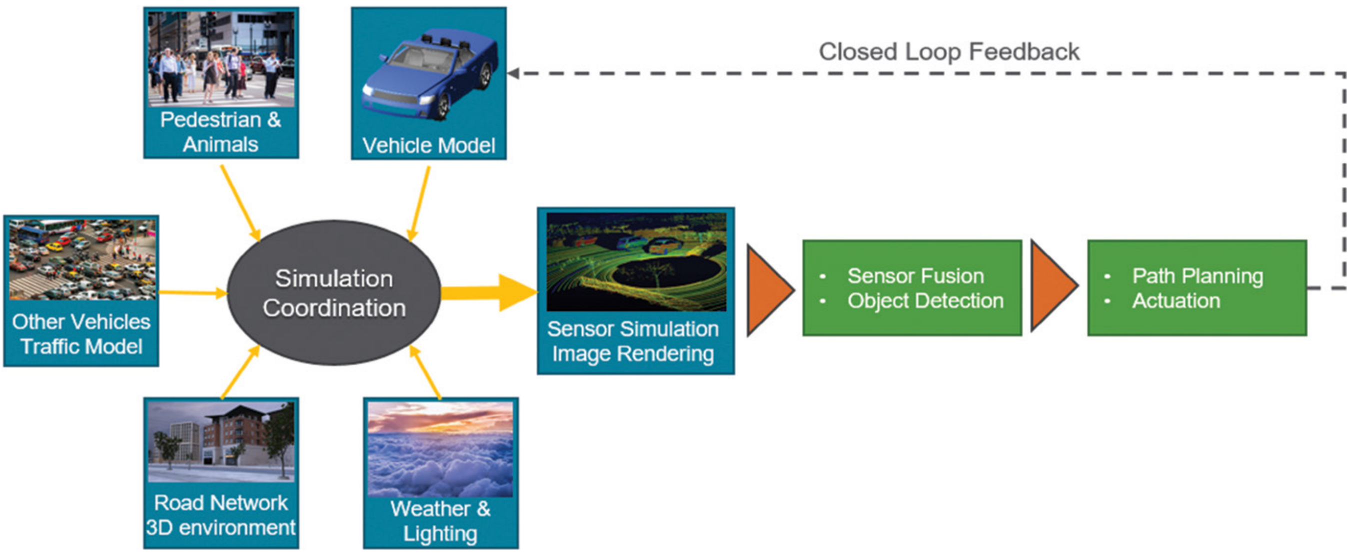

Figure 3. The end-to-end autonomous simulation workflow

autonomous vehicle is “as good as” a human driver, is 415 million kilometers. Do you claim your system is 20 times better? To statistically prove it, show me your results after 8 billion kilometers of testing!

By the way, every time someone modifies the position of a single sensor, the counter has to restart! Every time that some vehicle characteristic is changed (e.g. wheelbase, mass…), the counter has to restart. And every time a piece of software is updated… the counter has to restart!

Is there still someone that believes that road testing can bring you even close to full autonomy? Does anyone still believe that accumulating 1 million kilometers of road testing is a goal to be celebrated?

Take Waymo as an example (for the few of you who do not know them they are the “self-driving branch” of Google and they are considered the leaders in this space). Well, in 10 years they have accumulated 16 million kilometers of road data (really an outstanding outcome, even if most of these kilometers have been accumulated on sunny days in California and Arizona); at this rate, the richest company in the world would need over 200 years to prove they are “as good as human drivers”. And that’s why besides the road testing, Waymo is also running a fleet of 25 thousand virtual cars 24/7, simulating 13 million kilometers per day [6]. “Computer simulations are actually more valuable, as they allow manufacturers to test their software under far more conditions and stresses than could possibly be achieved on a test track.” said Ron Medford, Google’s safety director for the self-driving car program.

Everyone understands the necessity for road testing, but at the same time, we should notice that there are obvious drawbacks. Besides being slow and potentially dangerous if the testing is done prematurely, road testing is not repeatable or controllable, which are essential for autonomous system development.

To solve these issues, engineers tend to leverage the proving ground, which is much more repeatable. Moreover, real sensors can be evaluated on an actual vehicle. However, one of the disadvantages for the proving ground is the limited number of scenarios that engineers can test with. Each proving ground usually contains a set of scenarios and generally speaking it is slow and costly to build/construct new scenarios in the proving ground.

Now let’s take a look at simulation or virtual testing. In my opinion, there are a few key reasons why simulation is more applicable than road testing or proving ground for autonomous system development, especially in the initial phases of the project.

First, virtual testing is more scalable when it comes to cost. A fully equipped autonomous vehicle can cost up to half a million dollars, so a fleet of 200 vehicles would mean a 100 million dollars investment in the hardware itself (vehicles, sensors, data storage, wiring…). On the other hand, scaling virtual testing only requires you to have software licenses and CPU/GPUs to run the simulations, which is generally 100 times cheaper. Not to mention the operational cost to manage such a large fleet of vehicles (drivers, insurance, workshops, maintenance…). As an example of this scalability, BMW recently announced their new High-Performance-Cluster dedicated to the development of Autonomous Vehicles with more than 100,000 cores and more than 200 GPUs [7].

Secondly, the ground truth is always available in virtual testing. In the virtual environment, you always know if it is a pedestrian or if it is a car in front of you, and there is no need to hire service companies to do annotation/labeling for the road data that you collect from testing. When it comes to one billion miles of road data that is requested to validate the autonomous system, it is simply infeasible to annotate it all with human labor.

Thirdly, with simulation, engineers would be able to test the functions of the controller software in the early design stages. One would be able to test the different functions of the software separately with model-in-the-loop simulations without having to wait for the entire control system to be completed. Since you can replay the virtual scenarios as many times as you want, it’s much easier and cheaper to analyze, debug or iterate the core algorithms without having to consider the nuances of the actual production software.

“Many self-driving enthusiasts claim that ‘autonomous vehicles will reduce the number of people killed on the road by a factor of 20.’ But this a very bold statement that needs to be proven before being accepted by the general public, regardless of how appealing it sounds.”

Finally yet importantly, it is much more convenient to create permutations of a situation with virtual testing. Engineers can easily repeat the same test with a different set of parameters, like more pedestrians, higher speed, less sensor visibility, lower road friction, and many more. Permutations of a few basic scenarios with multiple parameters creates thousands of scenarios. And that's the key to ensure robustness and reliability of driving algorithms.

Autonomous Vehicle (AV) simulation is different from traditional vehicle simulations in a sense that apart from the vehicle itself, also the “environment” in which the vehicle operates is fundamental to assessing how it copes with all driving situations. The “environment” of an AV is quite rich (and sometimes crowded) as it includes all other vehicles, pedestrians, animals and of course the road, the sidewalks, buildings and even weather conditions. So, let's take a closer look into all these components.

To start with, the engineers need a vehicle model which represents the same dynamics characteristics as the actual vehicle. When you train the AI controller to drive the actual vehicle, the vehicle model needs to incorporate not only the correct mass and engine power, but also other correct behaviors like braking efficiency, or the load transfer during cornering events. All these performances are heavily influenced by the fundamental suspension designs (dampers, antiroll bars…) and the tire-road interactions.

Besides the vehicle model, the 3D environment also needs to be carefully constructed. 3D environments include the road network, which defines the space that the vehicle can occupy, and when and how the vehicle can occupy each lane. Besides the road itself, the immediate surroundings of the road is equally important. Trees and bushes can obstruct the view of the traffic signs, pedestrian from the sidewalks may suddenly decide to cross the street, buildings on the side of the street may cast shadows on to the road or reduce GPS accuracy. All these elements have to be realistically modeled to properly set the scene where the actions take place.

Of course, the autonomous vehicle shares the road with other vehicles, which can be bicycles, motorbikes, cars, buses, trucks with trailers, Segway, a police officer on a horse or anything else. Anything that is allowed to be driven on the road should be included in this case. And any of those participants might have their own way of interacting with the rest of the traffic. For example, a motorcycle splitting lanes during a traffic jam, while a large truck can easily get stuck in the traffic because of its slow acceleration, and a cyclist might decide to move from the sidewalk to the middle of the road to make a left turn. It is important that all those traffic participants be captured in their unique ways of maneuvering.

The pedestrian and their behaviors also need to be modeled, especially the way they interact with the oncoming traffic. The engineers need to reproduce the gestures of the pedestrians, for instance, whether or not they are watching the traffic, when they are distracted by texting on the phone while crossing the street. Animals' behaviors can be even more unpredictable, like jumping in front of the vehicle erratically, getting stuck in the middle of the road, or staring at the car when it's approaching.

The last important factor one needs to consider for the environment simulation is weather and lighting, which is critical since it impacts the way sensors perceive the scene. When it's raining outside, the vehicle needs to slow down because the driver's vision and the road friction have been changed. With the low-lying sun during sunset or sunrise, human driver needs to wear sunglasses because otherwise he/she couldn't really see the road clearly. Similarly, they also affect sensors like cameras, RADARs, or LiDARs. Fog reduces the visibility of a camera (and absorbs energy from a RADAR) and LiDARs are sensitive to rain drops since they scatter the laser beams.

Actually, the perceived sensor data is the most valuable piece of information that the AV simulation provides. With this data accurately available, engineers can focus on the following phases of the Autonomous Driving development.

Figure 4. Simulating pedestrian crossing the road while on cell phone. Simulation done in VIRES VTD.

Figure 5. Simulating vehicle driving during evening. Simulation done in VIRES VTD.

The first step is the so-called "sensor fusion" phase, in which data from different sensors is combined to calculate accurate position and orientation information. From the camera, the object is recognized, and when you associate the laser point cloud with the object, the distance to the object can be measured with LiDAR. And RADAR can even provide the speed of the object.

Now that the engineers have a clear understanding of the surroundings of the Ego vehicle, they can move on to the next phase, which is typically called "path planning". With the understanding of the pedestrians and the other vehicles in the traffic, the engineers need to predict what those other traffic participants will do in the next few tenths of second to the next few seconds. And essentially the vehicle needs to decide at that point what the safest thing is to do to cope with the situation.

Even with the sensor fusion, sometimes the situation is still not 100% understood by the AV (autonomous vehicle). If the vehicle is driving on the highway on a sunny day, all the sensors are giving the correct information and the vehicle has a clear view of a long distance ahead. But imagine if the autonomous vehicle is driving on a crowded street in New York during rush hour on a foggy day, you can't always tell if there're two pedestrians or three pedestrians in front of you. When the vehicle makes a decision as to which path to follow, it needs to consider not only the destination, but also what the safest route is to get there.

After the safest path is identified, it is time to decide on how to actuate the vehicle, which means how to apply the throttle, the brake, the steering wheel to follow that path, or how to adjust the damping in the suspension system to ensure a smooth ride. This is the so called "actuation phase" and is the playground of very specialized engineers that master control theory of ground vehicles.

In the virtual simulation workflow, all this information is being provided to the vehicle dynamics model as closed loop feedback. And based on those inputs with torques/forces, the vehicle model predicts its updated displacement, velocity and orientation to interact with the surroundings (including oncoming traffic or pedestrians that may be triggered to cross the street) and the simulation loop proceeds.

VIRES Virtual Test Drive (VTD) provides all the ingredients necessary for engineers to perform the autonomous driving simulation, and at the same time, VTD is compatible with not only MSC Software's internal technologies, but also with a number of 3rd party software. As an example of this, while VTD offers 2 different embedded techniques to represent the vehicle dynamics (with varying speed of simulation and accuracy of results) it can also be combined with Adams Car (the de facto standard in vehicle dynamics simulation) or with any other vehicle dynamics software. Same applies to the traffic models: VTD offers the most comprehensive traffic simulation capabilities in the industry (driving style of every vehicle can be set according to a number of parameters, and thousands of vehicle can be simultaneously simulated if necessary); nonetheless other traffic models can be incorporate into VTD as well, such as SUMO [8] or PTV Vissim [9].

Autonomous Driving is one of the most exciting, yet daunting tasks in the next decade to come. Road testing alone will never get us anywhere close to the billion miles of validation needed to ensure the safety of an autonomous car. In order to develop an autonomous driving system that can truly save tens of thousands of lives, comprehensive simulation of the real world is the key to success.

References

1. Washington Post https://www.washingtonpost.com/local/trafficandcommuting/deadly-driverless-uber-crash-spurs-debate-on-role-of-regulation/2018/03/23/2574b49a-2ed6-11e8-8688-e053ba58f1e4_story.html?utm_term=.d8bf81d9de6d
2. The Guardian https://www.theguardian.com/technology/2018/mar/22/self-driving-car-uber-death-woman-failure-fatal-crash-arizona
3. NHTSA https://crashstats.nhtsa.dot.gov/Api/Public/ViewPublication/812603
4. Rand Corporation https://www.rand.org/content/dam/rand/pubs/research_reports/RR1400/RR1478/RAND_RR1478.pdf
5. Assuming that about 95% of road deaths are caused by human error
6. https://waymo.com/safety/
7. https://www.press.bmwgroup.com/global/article/detail/T0293764EN/the-new-bmw-group-high-performance-d3-platform-data-driven-development-for-autonomous-driving?language=en
8. SUMO is an open source traffic model originally create by DLR. For more information, please see https://sumo.dlr.de/index.html
9. This functionality will be added in VTD 2019.1 http://vision-traffic.ptvgroup.com/en-us/products/ptv-vissim/