December 31, 2019
California Department of Motor Vehicles
Autonomous Vehicle Program
ATTN: Mr. Miguel Acosta, Chief Autonomous Vehicles Branch
P.O. Box 932342
MS L224
Sacramento, CA 94232-3420

Dear Mr. Acosta:

Aurora's mission is to deliver the benefits of self-driving technology safely, quickly, and broadly. Our team has decades of experience developing self-driving technology, and has seen many different approaches taken to its testing. In the early days of this technology, on-road autonomous driving was the only tool the industry had for surfacing issues with the system's hardware and software.

This is no longer true.

While on-road testing is useful for collecting targeted data and performing late stage validation of self-driving systems, we find that large-scale, on-road autonomous testing is a slow, and inefficient approach to development relative to more sophisticated, virtual techniques. So instead of prioritizing the collection of on-road miles in 2019, we've continued the trajectory we started when we founded the company, using on-road testing for targeted data collection in increasingly complex environments and investing heavily in a robust suite of virtual testing tools that runs repeatedly and reliably at scale on a growing body of real and simulated data. In 2019, we've paired these improvements with several investments into the Aurora Driver hardware and infrastructure across the company. Notably, this year saw us integrate the Aurora Driver with three new vehicle platforms, perform two major hardware upgrades that significantly increased both computational power and sensor capability, and make countless improvements to infrastructure across the company. These include expanding our HD mapping system (the Aurora Atlas), strengthening our machine learning pipelines, and building custom tools that allow us to extract meaningful metrics from many forms of data.

We've been pleased with the results. In 2019, our virtual testing suite has grown to over 735k tests per day, an increase of over 100x from 2018. As this suite has grown, our targeted data collections have pushed more deeply into complex urban areas. In a typical mile of driving in 2019, our vehicles encounter approximately 3x more vehicles and nearly 10x more pedestrians than they did in 2018. This has allowed us to keep our testing fleet small and drive fewer miles on public roads as these vehicles are both collecting more interesting data and, through our virtual test suite, we are doing more with it.

Thanks to our growing virtual testing capabilities, and despite this increased operational complexity, our vehicles now perform far more complex maneuvers than they did in 2018. Our fleet of test vehicles now yields to pedestrians entering or exiting parked vehicles, jaywalking individually or in groups, or crossing at crosswalks. It regularly navigates complex, unprotected left turns or rights on red, reasons about gridlock, intersection ownership, right of way, and movements of various classes of actors, and safely maneuvers around double-parked vehicles or other objects in its lane.

Finally, and though they make up a smaller portion of our approach, on-road miles still serve three specific purposes. The first is to evaluate if our system behaves as expected in various environmental circumstances, such as weather, dust and pollen, glare, and smog. The second is to inspire new and more lifelike offline testing scenarios that we then use to validate our system. The third purpose of on-road miles is to collect training data that strengthens the machine learning models throughout our software system, most of which can be done more efficiently and effectively with manual driving. In 2019, we drove 39,729 miles on California public roads, of which 26,300 miles were manually driven. While testing in autonomous mode, our vehicles experienced 142 reportable disengagements, approximately 25% of which were caused by a single software issue discovered and addressed early in the year. Despite this, our rate of disengagements per 1,000 autonomous miles is similar this year (10.6) to what it was in last year's report (11.5). This is also by design, and is consistent with the philosophy we outlined in 2018.

Overall, it's been a great year for the development of the Aurora Driver. For the first time, we were able to demonstrate our initial thesis for development: an expansive suite of carefully crafted virtual tools allows quicker development than on-road driving. Our engineers make more progress more quickly today than they did a year ago, despite driving fewer miles on California roads. While we recognize that on-road testing miles are an essential tool for the ultimate safety and deployment of self-driving vehicles, we are excited about the virtual testing infrastructure we've created, the driving capabilities we've built, and the trajectory we're on.

Sincerely,

Chris Urmson
Co-Founder and CEO, Aurora

Appendix A: Disengagement Details

The table summarizes all critical disengagements by date, vehicle, initiator, location, and cause (OL 311R Section 2). The causes can be one of the following:

Perception discrepancy refers to scenarios in which the perception system fails to detect an object correctly. An example of a perception discrepancy would be when the self-driving system inaccurately otherwise failed to track the vehicle ahead of it.

Planning discrepancy refers to cases where, in the judgment of our trained vehicle operators, the planned path of the vehicle was unacceptable, unwarranted, or unexpected given the driving scenario encountered. An example of a planning discrepancy would be when the self-driving system provided a poor trajectory for the vehicle to change lanes.

Control discrepancy refers to scenarios in which the proposed or actual vehicle control created an undesired behavior. An example of a control discrepancy would be when, in the judgement of our trained vehicle operators, the vehicle responded to a given driving scenario with excessive braking.

Localization discrepancy refers to scenarios in which the autonomous system was unable to confidently ascertain its location in the world and gave control back to the human operator.

All vehicles listed are not yet capable of operating without a driver (though that is our ultimate intent) and all had a safety driver and copilot present.