Transportation Research Part A 94 (2016) 182–193

Contents lists available at ScienceDirect

Transportation Research Part A

journal homepage: www.elsevier.com/locate/tra

Driving to safety: How many miles of driving would it take to demonstrate autonomous vehicle reliability?

Nidhi Kalra *, Susan M. Paddock

RAND Corporation, 1776 Main Street, Santa Monica, CA 90401, United States

ARTICLE INFO

Article history:
Received 22 March 2016
Received in revised form 19 July 2016
Accepted 6 September 2016

Keywords:
Autonomous vehicles
Vehicle automation
Safety
Reliability
Test driving
Statistical significance

ABSTRACT

How safe are autonomous vehicles? The answer is critical for determining how autonomous vehicles may shape motor vehicle safety and public health, and for developing sound policies to govern their deployment. One proposed way to assess safety is to test drive autonomous vehicles in real traffic, observe their performance, and make statistical comparisons to human driver performance. This approach is logical, but it is practical? In this paper, we calculate the number of miles of driving that would be needed to provide clear statistical evidence of autonomous vehicle safety. Given that current traffic fatalities and injuries are rare events compared to vehicle miles traveled, we show that fully autonomous vehicles would have to be driven hundreds of millions of miles and sometimes hundreds of billions of miles to demonstrate their reliability in terms of fatalities and injuries. Under even aggressive testing assumptions, existing fleets would take tens and sometimes hundreds of years to drive these miles—an impossible proposition if the aim is to demonstrate their performance prior to releasing them on the roads for consumer use. These findings demonstrate that developers of this technology and third-party testers cannot simply drive their way to safety. Instead, they will need to develop innovative methods of demonstrating safety and reliability. And yet, the possibility remains that it will not be possible to establish with certainty the safety of autonomous vehicles. Uncertainty will remain. Therefore, it is imperative that autonomous vehicle regulations are adaptive—designed from the outset to evolve with the technology so that society can better harness the benefits and manage the risks of these rapidly evolving and potentially transformative technologies.

1. Introduction

In the United States, roughly 32,000 people are killed and more than two million injured in crashes every year (Bureau of Transportation Statistics, 2015). U.S. motor vehicle crashes as a whole can pose economic and social costs of more than $800 billion in a single year (Blincoe et al., 2015). And, more than 90 percent of crashes are caused by human errors (National Highway Traffic Safety Administration, 2015)—such as driving too fast and misjudging other drivers' behaviors, as well as alcohol impairment, distraction, and fatigue.

Autonomous vehicles have the potential to significantly mitigate this public health crisis by eliminating many of the mistakes that human drivers routinely make (Anderson et al., 2016; Fagnant and Kockelman, 2015). To begin with,

* Corresponding author.
E-mail addresses: nidhi_kalra@rand.org (N. Kalra), susan_paddock@rand.org (S.M. Paddock).

http://dx.doi.org/10.1016/j.tra.2016.09.010
0965-8564/

autonomous vehicles are never drunk, distracted, or tired; these factors are involved in 41 percent, 10 percent, and 2.5 percent of all fatal crashes, respectively (National Highway Traffic Safety Administration, 2011; Bureau of Transportation Statistics, 2014b; U.S. Department of Transportation, 2015).[1] Their performance may also be better than human drivers because of better perception (e.g., no blind spots), better decisionmaking (e.g., more-accurate planning of complex driving maneuvers like parallel parking), and better execution (e.g., faster and more-precise control of steering, brakes, and acceleration).

However, autonomous vehicles might not eliminate all crashes. For instance, inclement weather and complex driving environments pose challenges for autonomous vehicles, as well as for human drivers, and autonomous vehicles might perform worse than human drivers in some cases (Gomes, 2014). There is also the potential for autonomous vehicles to pose new and serious crash risks, e.g., crashes resulting from cyber-attacks (Anderson et al., 2016). Clearly, autonomous vehicles present both enormous potential benefits and enormous potential risks.

Given the high stakes, policymakers, the transportation industry, and the public are grappling with a critical concern: How safe should autonomous vehicles be before they are allowed on the road for consumer use? For the answer to be meaningful, however, one must also be able to address a second concern: How safe are autonomous vehicles?

Perhaps the most logical way to assess safety is to test-drive autonomous vehicles in real traffic and observe their performance. Developers of autonomous vehicles rely upon this approach to evaluate and improve their systems,[2] almost always with trained operators behind the wheel who are ready to take control in the event of an impending failure incident.[3] They can analyze the failure incident after the fact to assess what the autonomous vehicle would have done without intervention, and whether it would have resulted in a crash or other safety issue (Google, 2015). Developers have presented data from test driving to Congress in hearings about autonomous vehicle regulation (Urmson, 2016).

But is it practical to assess autonomous vehicle safety through test-driving? The safety of human drivers is a critical benchmark against which to compare the safety of autonomous vehicles. And, even though the number of crashes, injuries, and fatalities from human drivers is high, the rate of these failures is low in comparison with the number of miles that people drive. Americans drive nearly 3 trillion miles every year (Bureau of Transportation Statistics, 2015). The 2.3 million reported injuries in 2013 correspond to a failure rate of 77 reported injuries per 100 million miles. The 32,719 fatalities in 2013 correspond to a failure rate of 1.09 fatalities per 100 million miles.

For comparison, Google's autonomous vehicle fleet, which currently has 55 vehicles, was test-driven approximately 1.3 million miles in autonomous mode and was involved in 11 crashes from 2009 to 2015.[4] Blanco et al. (2016) recently compared Google's fleet performance with human-driven performance. They found that Google's fleet might result in fewer crashes with only property damage, but they could not draw conclusions about the relative performance in terms of two critical metrics: injuries and fatalities. Given the rate of human and autonomous vehicle failures, there were simply not enough autonomously driven miles to make statistically significant comparisons.

In this report, we answer the next logical question: How many miles[5] would be enough? In particular, we first ask:

1. How many miles would autonomous vehicles have to be driven without failure to demonstrate that their failure rate is below some benchmark? This provides a lower bound on the miles that are needed.

However, autonomous vehicles will not be perfect and failures will occur. Given imperfect performance, we next ask:

2. How many miles would autonomous vehicles have to be driven to demonstrate their failure rate to a particular degree of precision?
3. How many miles would autonomous vehicles have to be driven to demonstrate that their failure rate is statistically significantly lower than the human driver failure rate?

[1] This does not mean that 53.5 percent of all fatal crashes are caused by these factors because a crash may involve, but not be strictly caused by, one of these factors, and because more than one of these factors may be involved in a single crash.

[2] Extensive testing on public roads is essential for developing and evaluating autonomous vehicles, given their great complexity and the diversity and unpredictability of conditions in which they need to operate. In contrast, typical automobile components are significantly simpler and their operating conditions can be well defined and recreated in controlled settings, which enables laboratory testing and verification. Curtain-style air bags, for example, are tested with a combination of component tests to assess inflation time, fill capacity, and other responses in a range of temperature conditions and impact configurations, as well as laboratory crash testing to evaluate their performance in collisions (Kaleto et al., 2001).

[3] Some states, such as California, require trained drivers to be behind the wheel of any autonomous vehicle driving on public roads (California Vehicle Code, 2012).

[4] Two of these crashes involved injury and none involved a fatality. Seven of the crashes did not reach a level of severity that would warrant a Department of Motor Vehicles report (Blanco et al., 2016).

[5] Note that not all miles of road are created equal. The miles used to demonstrate autonomous vehicle safety must represent the full range of conditions (climate, terrain, congestion, etc.) in which humans drive, and be proportionally distributed as well. That is, if 10 percent of human-driven miles occur in snow, so too must the autonomous vehicle test miles.