Autonomous Vehicle Safety and Deployment: Lessons from Human Crashes

Akhil Shetty, Hamidreza Tavafoghi, Alex Kurzhanskiy,
Kameshwar Poolla, and Pravin Varaiya

University of California, Berkeley

Abstract

Traffic crashes claim more than 36,000 lives every year in the US, with economic costs amounting to $836B. Autonomous vehicles (AVs) promise to usher in a future with near-zero crashes. It seems likely that widespread deployment of AVs will eliminate the large number of crashes caused by impaired, distracted or reckless drivers. However, it is unclear whether AVs will be able to avoid a significant fraction of the remaining crashes for which no driver is directly responsible. As such, deploying AVs without adequately assessing their safety capabilities might lead to an increase in crashes rather than a reduction. In this paper, we discuss how an analysis of human crashes can provide insights about the types of crashes that remain challenging for AVs and the role of connected infrastructure in addressing them. We also discuss how human crashes and driving data can be valuable in inferring safety capabilities of AVs in diverse driving contexts. Based on these observations, we provide suggestions for policies and regulations governing the deployment of AVs.

1 Introduction

Traffic crashes claim more than 36,000 lives annually in the US [1]. About 20% of these fatalities are vulnerable road users such as pedestrians and bicyclists. Cumulatively, traffic crashes result in $836B in comprehensive economic costs every year [2]. Autonomous vehicles (AVs) promise to lead us to a safe transportation future by eliminating error-prone humans from the driver's seat. AV companies claim that their technology, once fully developed, will eliminate 94% of all crashes that are attributed to human error [3]. The promise of massive safety gains has helped the AV industry attract more than $100B in investment over the last decade [4]. Initially, there was widespread optimism that AVs would be widely available by 2020.

However, the vision of widespread deployment of AVs has not yet materialized. While AV testing and deployment has commenced in limited regions, there are several safety challenges that need to be addressed for mass adoption of this technology. AV companies expect to eliminate most crashes as their technology improves over time. Recent testing offers promising signs about the ability of AVs to prevent a significant fraction of fatal crashes [5]. At the same time, fatal crashes involving AVs show that faulty technology can cost lives rather than save them [6, 7].

The often quoted 94%-statistic from NHTSA [3] seems to suggest that most crashes are caused by human error and hence, AVs can eliminate these crashes simply by avoiding errors. However, a closer reading of the NHTSA report calls this narrative into question. The report explicitly clarifies that the attribution of 94% of crashes to human error is not meant to suggest that drivers are at fault in all these cases. As pointed out in [8], a significant fraction of these crashes are a consequence of the difficulty of driving rather than an obvious mistake made by the driver. Several studies predict that the actual crash reduction brought about by AVs will be closer to 50% rather than 94% [8, 9, 10].

Since AVs will be sharing the road with other human drivers, pedestrians and bicyclists, errors made by them can prove to be fatal. Thus, AV companies need to demonstrate their safety before they can be widely deployed. This is a challenging task due to the diverse range of situations encountered while driving.

To be deemed safe, AVs must prove that they can navigate through a wide variety of environments, road user behavior and weather conditions. During on-road testing, AVs encounter only a tiny fraction of the overwhelming diversity of contexts that human drivers come across during billions of miles driving every year. An analysis of human crashes can point us to particularly challenging driving contexts that AVs might not have faced during testing. Moreover, since some crashes are likely to occur even with improved AV technology, human crash rates across diverse driving contexts can act as a useful benchmark to interpret the safety capabilities of AVs.

Policy makers and regulators are faced with the challenge of balancing the safety benefits of AV technology and its potential pitfalls. AVs have the potential to save thousands of lives every year. However, allowing the deployment of AVs with faulty technology can end up costing lives. Regulators are tasked with deciding whether AVs are safe enough to be deployed on the roads. Merely analyzing the number of crashes or disengagements AVs were involved in during testing is insufficient to infer their safety prowess. For instance, an AV that has driven crash-free on sparse highways may not necessarily be safe enough to navigate through dense, urban intersections. Moreover, some crashes can occur despite AVs not making errors. In order to assess the safety capabilities of AVs, it is important to have a clear understanding about the types of crashes AVs can or cannot avoid based on current technology. Expecting AVs to never be involved in crashes would be unrealistic. At the same time, it is reasonable to ensure that AVs do not make sensing errors similar to those which led to Uber's fatal crash in Tempe, Arizona [6]. Human crash data can help regulators come up with acceptable crash risk baselines that AVs should satisfy to be allowed to operate in their jurisdictions. It can also help inform policies that steer the development of AV technology towards preventing crashes that are unavoidable at present.

In this paper, we discuss how an analysis of human crashes can provide insights about the types of crashes that remain challenging for AVs. Recognizing that some crashes are inevitable, we discuss how human crashes and driving data can be used to infer safety capabilities of AVs in diverse driving contexts. Based on these observations, we provide suggestions for policies and regulations governing the deployment of AVs.

2 Understanding Crash Causes

In order to understand the types of crashes AVs can or cannot avoid, we must first understand why they occur. Due to the complex nature of driving, a variety of circumstances can lead to crashes. In addition, crashes are rare events. In the US, a crash occurs once every 500,000 miles on average [11]. Moreover, a fatal crash occurs once every 100 million miles. Since AV fleets have only been tested over tens of millions of miles in limited operational design domains (ODDs), merely analyzing testing data is insufficient for understanding which crashes they can avoid. For instance, Waymo AVs have covered about 20 million miles in on-road testing since 2009 [12]. On the other hand, humans drive billions of miles every year encountering a considerably wider variety of road geometries, weather conditions and neighboring vehicle behaviors than the AVs being tested at present. As a result, they come across a significantly greater diversity of crash modes than their AV counterparts. Analyzing human crash data can provide a comprehensive picture of crash causes. NHTSA periodically provides a crash cause based breakdown of a representative sample of national crashes [11, 13]. In order to understand which of these crashes AVs can successfully avoid, we compile the following taxonomy.[1]

2.1 A Crash Taxonomy

As shown in Figure 1, crashes are categorized at the first level based on whether they were attributed to driver-related causes. Non-driver related causes included vehicle failures (eg., faulty tires, failed brakes) and challenging environmental conditions (eg., slick roads, fog, bad road design). These non-driver related causes account for 4% of all crashes. Since these crashes were not caused by decisions made by the driver, AVs are unlikely to avoid such crashes at present. Nevertheless, it is conceivable that better automotive technology

[1]The crash cause could not be identified for 2% of all crashes in studied in [13]. This explains why driver and non-driver related crashes add up to 98% of all crashes.

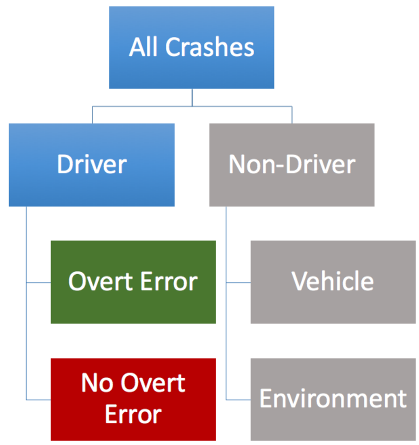

Figure 1: A taxonomy of crashes based on crash causes.

will eliminate vehicle failure crashes over time. Some environment-related crashes could also be prevented by more cautious driving and better road design.

Driver-related causes account for 94% of all crashes. This statistic is often quoted by AV companies to point to the potential safety benefits that can be brought about by AVs. However, not all driver-related crashes are *caused* by drivers. The NHTSA report explicitly states the following: "Although the critical reason is an important part of the description of events leading up to the crash, it is not intended to be interpreted as the cause of the crash nor as the assignment of the fault to the driver, vehicle, or environment." Thus, drivers are not necessarily at fault for 94% of all crashes. As such, it is not clear whether AVs can prevent all of these crashes. This motivates a more refined classification of driver-related crashes based on whether they involved drivers making overt errors.

2.2 Safety Opportunities and Challenges

Overt driver errors include impaired driving, speeding and violating traffic rules. Impaired and distracted driving are involved in 28% and 9% of all fatal crashes respectively [11]. Traffic violators cause 15% of fatal crashes. It seems reasonable to expect that AVs will eventually eliminate such crashes, saving close to 20,000 lives that are lost to these causes every year. This clearly indicates the immense potential AVs possess to reduce traffic fatalities. Several studies suggest that the fraction of all crashes caused by overt errors is closer to 50% as opposed to 94% [8, 9, 10].

Nevertheless, this implies that a significant fraction of crashes occur despite drivers not making overt errors. An example of this is crashes involving occlusions on the road which prevent even attentive drivers from detecting each other. These crashes are a consequence of the difficulty of driving among other road users with conflicting trajectories, partially observed positions and unpredictable behavior. Scenarios involving such challenges are routinely observed particularly in dense, urban settings. AVs have been known to struggle with maneuvers such as left turns and merging that involve predicting the positions and behavior of other vehicles and pedestrians [14, 15, 16]. This suggests that AVs may not be able to avoid all crashes involving gaps in information about the positions and actions of neighboring road users. We will explore such *information gaps* in greater detail in Section 4 and discuss how AVs can deal with them.

Although AVs will not drive impaired or distracted, they consist of hardware and software subsystems that can be faulty and hence, create new crashes in addition to those involving information gaps. The fatal crash involving an Uber AV colliding with a pedestrian in Tempe, Arizona is an important example to keep in mind. In this incident, the AV's Lidar alternated between misclassifying the pedestrian as "vehicle", "unknown object", and "bicycle" in the seconds leading up to the crash [17]. In order to ensure that such errors do not claim more lives, regulators need to ensure that only demonstrably safe AVs are permitted to be deployed in their jurisdictions. In the next section, we discuss how human crashes can provide informative signals in the risk assessment of AVs.

3 Risk Assessment

While AVs can eliminate a significant fraction of prevailing crashes, they can also create new ones. Introducing AVs with faulty technology to the roads can cost lives [6, 7]. Thus, risk assessment of AVs is crucial before they are deployed on a large scale. In the last few years, several AV companies have been testing their fleets in limited ODDs. Presumably, such testing can provide insights about the safety capabilities of AVs. However, there is no standard approach to assess AV performance at present. This raises the question: *How should AV safety capabilities be inferred from on-road testing?*

3.1 Diversity of Driving Contexts

The crashes per mile metric has been traditionally used to describe human safety performance. The human crash rate is 1.9 per million miles[2]. As AV fleets cover tens of millions of miles every year, a statistical comparison between the human and AV crash rate becomes more plausible. However, such a comparison might result in flawed conclusions. Humans drive billions of miles every year across a wide variety of road geometries, weather conditions and driving cultures. On the other hand, AVs have only been tested over millions of miles in very limited domains in which they were deemed safe enough to operate. As we saw in the previous section, there are certain road scenarios that are considerably more dangerous than the rest. The crashes per mile metric abstracts away the diversity of contexts encountered while driving and hence, is insufficient to make conclusions about safety performance.

In order to derive insights about AV safety capabilities, we need to know not just the number of miles, but also where they were driven and what maneuvers were involved. For instance, testing over millions of miles on sparse highways suggests very little about their ability to make a left turn at a busy intersection.

3.2 Incorporating Diversity in Risk Assessment

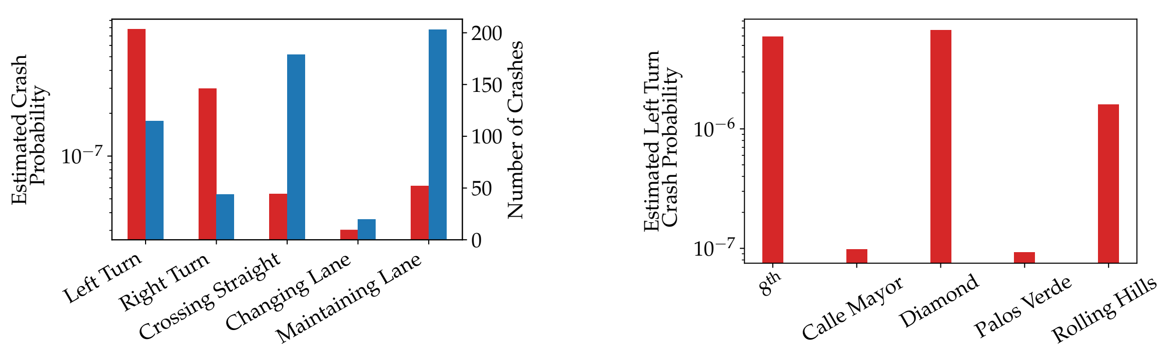

Figure 2: Estimated crash probabilities along the Pacific Coast Highway in Torrance, California (Adapted from [18]). (Left): Estimated crash probabilities (red) and number of crashes (blue) across diverse maneuvers. (Right): Estimated left turn crash probabilities at various intersections.

There are two major sources of diversity in driving: maneuvers (eg., left turn, changing lanes) and environments (eg., location, weather, time of day). In [18], we developed a risk assessment framework that takes into account the diversity of maneuvers and environments. An analysis of human crash risk across diverse driving contexts can provide broader context to the observations made during on-road testing of AVs. Using crash reports and traffic data, we derived crash risk estimates for an average human driver over a 12-mile stretch of the Pacific Coast Highway in Torrance, California. We found that crash probabilities

[2] This is based on police-reported crashes. A NHTSA report [2] estimates that 40% of all crashes go unreported to the police. Taking this into account, the actual human crash rate turns out to be 3.2 per million miles.

vary considerably based on maneuver (see Figure 2). The probability of a crash during a left turn was estimated to be 7.8×10^{-7}, while that for a lane change turned out to be 3×10^{-8} . A crash during a left turn was 26 times as likely as that during a lane change. Moreover, even for the same maneuver, crash probabilities varied significantly based on location. The left turn crash probability for the most dangerous intersection on the above stretch was two orders of magnitude higher than that for the least dangerous one.

While the above context-aware analysis was performed for human drivers, it can provide refined insights about challenges faced by AVs across diverse driving contexts. As a case in point, a recent safety report by Waymo provides a maneuver-based breakdown of crashes involving its AVs in Maricopa County, Arizona [12]. The report suggests that Waymo AVs were able to avoid road departure related crashes over 6 million miles of driving. This is a positive sign of their ability to avoid such crashes that are usually a result of aggressive or impaired driving, accounting for 27% of all fatal crashes in Maricopa County. On the other hand, they were involved in 47 crashes involving maneuvers such as lane changes, left/right turns and crossing intersections. Since data about the exact locations traversed and the frequency of various maneuvers performed by the AVs is unavailable, only limited insights can be derived about their safety capabilities in such driving contexts. Thus, in addition to a detailed breakdown of crash causes, data disclosure about the maneuvers and environments in which they were performed is required to make an informed assessment of AV safety capabilities. Combining a comprehensive summary of AV on-road testing with a context-aware analysis of human crashes would enable us to: (i) interpret the on-road performance of AVs compared to human drivers in similar contexts, and (ii) forecast how their safety performance will generalize to other regions where they haven't been tested. We will discuss in Section 5 how such a risk assessment framework can inform the decision-making process of regulators regarding the deployment of AVs in their jurisdictions.

4 Information Gaps and Connectivity

Being safe on the roads depends not just on our own actions but also on the positions and actions of neighboring vehicles, both of which are only partly known to us. Gaps in such safety critical information can result in crashes despite vehicles not making an obvious error. We illustrate below how *information gaps* can cause crashes through a common on-road scenario.

4.1 Information Gaps: An Illustrative Example

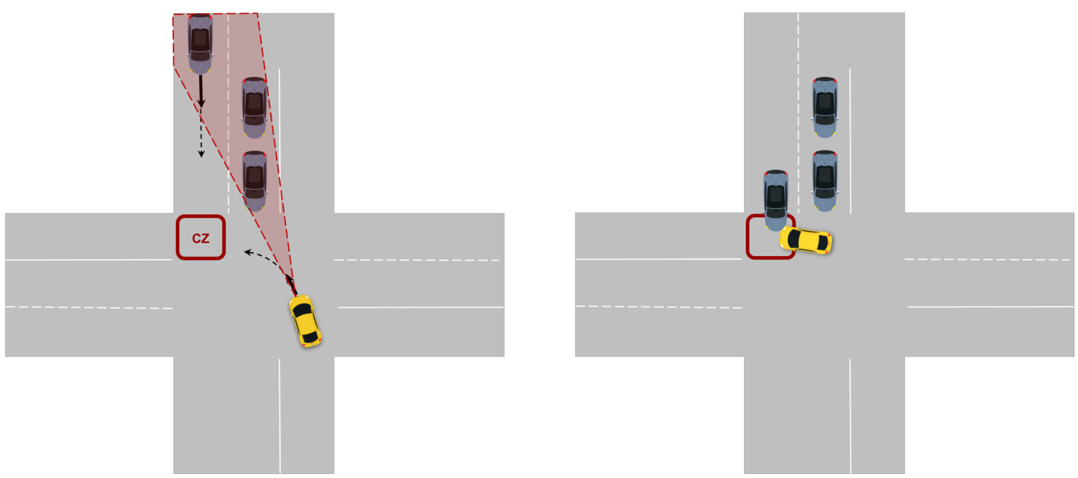

Figure 3: (Left): Yellow vehicle making an unprotected left turn with through moving vehicles occluded by queued up vehicles in the middle lane. The left turning vehicle's occluded field of view is illustrated by the shaded red region. The red box *CZ* denotes the conflict zone. (Right): The left turning and through moving vehicles do not detect each other due to the occlusion and hence, cannot stop in time to avoid a crash (adapted from [19]).

Consider the scenario depicted in Figure 3 where the yellow vehicle is attempting an unprotected left turn. The vehicles in the opposing through-moving lane (left-most lane in Figure 3) have right-of-way and hence, the yellow vehicle must yield to them. The vehicles in the middle lane are queued up, waiting for an opportunity to make a left turn. The presence of these vehicles occludes the field of view of the left-turning and through-moving vehicles. As a result, both vehicles keep moving until they see each other, at which point, it is too late to avoid a collision.

This example illustrates how the lack of information about the positions of vehicles with conflicting paths can result in a crash despite the drivers not making an obvious error. In fact, an Uber AV was involved in a similar crash during on-road testing in March 2017 [20]. More generally, knowledge about both the position as well as behavior of neighboring vehicles is a pre-requisite to be safe on the roads. However, we routinely encounter gaps in this knowledge in a variety of driving scenarios. Merging on to a freeway is a scenario which requires a driver to predict how neighboring vehicles will respond to their actions. Red-light violations are an example of our assumption about the rule-following behavior of neighboring vehicles breaking down, resulting in hazardous scenarios. Since safety in these scenarios depends on not just our own behavior but also the position and behavior of other vehicles, it is not immediately clear whether better AV technology can prevent such crashes.

The predominant approach in the AV industry is to divide the complex task of driving into sub-tasks such as sensing, perception, prediction, planning and control. The expectation is that with improvements in technology, AVs will get better at performing each of these sub-tasks and eventually eliminate most crashes. However, the lack of connectivity in current AV designs poses a major safety challenge especially in driving scenarios involving information gaps arising from occluded view or unpredictable behavior of other road users. A key question then is: *In the absence of connectivity, can AVs be safe despite information gaps?*

4.2 Safety in the Absence of Connectivity

Connectivity with neighboring vehicles or infrastructure (V2V/I2V) would eliminate such information gaps [21, 22, 23]. In the example above, a sensor placed on the through-moving lane along with I2V communication would ensure that the vehicles detect each other in time to avoid a crash. However, most AV companies do not rely on communication with neighboring vehicles or infrastructure to ensure safety. For an AV to be safe despite information gaps, it must be robust to all positions and actions of neighboring vehicles that are consistent with its partial observations. Such a worst-case safe approach was first proposed by Mobileye through its Responsibility Sensitive Safety framework in [24]. Similar proposals have been subsequently proposed by other companies for navigating such situations [25, 26]. In [19], we show that following such worst-case safe approaches would preclude AVs from making maneuvers such as unprotected left turns and merging in common on-road settings. More importantly, even if AVs were to tolerate a crash risk comparable to human drivers, we find that they would need to act more conservatively in order to resolve such information gaps. For the above unprotected left turn scenario, our study suggests that an AV would need to wait for more than 3 minutes at the intersection to make sure that it achieves a crash risk comparable to humans. Thus, forgoing connectivity indeed comes at a significant cost to either safety or throughput.

5 Deploying Automated Vehicles: Policies and Regulation

As the deployment of AVs expands over time, policy makers and regulators need to grapple with the challenge of balancing the risks associated with this burgeoning technology and its potential to make our roads safer. While AVs can eliminate prevailing crashes involving overt human error such as impaired or reckless driving, faulty technology can cause crashes that wouldn't have occurred otherwise. Since errors made by AVs can result in loss of life, regulators need to make sure that all the necessary checks are carried out before allowing deployment. As a case in point, the National Transportation Safety Board concluded that the absence of regulatory oversight contributed to the fatal crash involving an Uber AV in 2018 [17]. At the same time, certifying that AVs have a fatality rate significantly lower than the average human driver would require billions of miles of on-road testing [27]. However, delaying the introduction of AVs until they are certified safe would result in thousands of traffic fatalities that could have been potentially saved with this technology [28]. We now discuss how the lessons learnt from previous sections can inform policies and regulations governing AV deployment.

5.1 Testing Data Disclosure

AV companies collect vast amounts of data during on-road testing. Analyzing this data can provide insights about the current safety capabilities of AVs, which in turn can inform policies regarding their deployment. Regulations regarding data disclosure vary significantly across states. While states like Arizona and Florida do not require any data disclosure, California requires AV companies to disclose every disengagement their fleets were involved in. Most companies are opposed to this requirement [29, 30]. Note that apart from disengagements, companies are only required to report the number of miles driven by their fleets. One cause of concern is that this leads to a comparison of disengagements per mile across companies without any context about the environments in which these miles were driven. This creates perverse incentives for AV companies involved in testing. For instance, companies can test their fleets only on highways in rural California during off-peak hours in order to report a lower disengagement rate compared to other companies that test their AVs in the busy streets of San Francisco. Moreover, what constitutes a disengagement is at the discretion of the company. A focus on disengagement rates incentivizes companies to increase the danger threshold at which they declare a disengagement and hand off control to the safety driver [31]. This jeopardizes the safety of other road users.

Thus, regulators need to take into account the diversity and frequency of contexts encountered by AVs to have a well-informed opinion about their safety capabilities. As discussed in Section 3, one approach along these lines is to have companies report the frequency of maneuvers and the environments in which they were performed. This allows an apples-to-apples comparison between the performance of various AV companies. Additionally, such a summary of AV safety performance can be complemented with a context-aware analysis of human crashes. This would enable an informed assessment of AV safety capabilities in various driving contexts and how they compare with human drivers. Moreover, it allows us to forecast how AV performance in specific ODDs will translate to jurisdictions where they haven't yet been deployed.

5.2 Permitting AV Deployment

In recent years, several jurisdictions have allowed the deployment of AVs in specific ODDs [32]. As a case in point, Waymo currently operates a driverless ride-hailing service within a 100-square-mile ODD in Phoenix, Arizona. Since AVs will be driving without safety drivers in such regions, faulty technology can prove to be fatal. Moreover, the negative public sentiment surrounding AVs as a result of Uber's fatal crash in 2019 indicates that AV deployment can suffer a major setback if such incidents are repeated. Thus, regulators need to ensure that due diligence is done before they permit AV companies to operate on their roads.

Regulators can gain insights about the safety capabilities of AVs by analyzing on-road testing data. This would inform them about how effective AVs are in safely navigating through a diverse set of traffic scenarios such as left turns at a busy intersection or changing lanes with narrow vehicle gaps. An analogy with human driving tests is instructive here. During a driving test, drivers are required to demonstrate their ability to perform a diverse set of maneuvers such as parking, lane changes and turns. As discussed in Section 3, there is substantial variation in crash risk across environments even for the same maneuver. Thus, a "driving test" for AVs should involve an assessment of their ability to perform a host of maneuvers at locations known to be hazardous for humans. Such an assessment can be done so long as adequate context is provided in the testing data submitted to regulators, as described in 5.1. Allowing AVs to be tested and deployed without a safety driver necessitates greater caution. The rarity of crashes in terms of vehicle miles travelled suggests that thousands of miles of testing in a particular region is insufficient to be completely assured about an AV's safety prowess. A risk mitigating strategy here is to undertake a phased deployment of AVs from "easy" to "difficult" regions in the jurisdiction. This will make sure that obvious errors in the technology pipeline are caught early on without inflicting too much damage to other road users. As an example, intersection-related maneuvers are known to cause difficulties for AVs [14, 15]. Thus, it makes sense to start deployment on highways during off-peak hours and gradually move towards allowing AVs to operate near busy intersections.

5.3 Investment in Connected Infrastructure

We saw in Section 4 how information gaps can cause crashes in the absence of connectivity. However, there are several disincentives to AV companies adopting I2V/V2V connectivity at present. I2V requires instrumentation of infrastructure that can be expensive. Both I2V and V2V requires vehicles to have the required communication apparatus installed which further adds to the cost. Moreover, relying on such connectivity limits the deployment of AVs to regions in which the penetration rate of connected vehicles and infrastructure is high. Since this is true for only a few regions as of now, it severely constrains the region of operation of AVs, resulting in a weak business case. Moreover, there are cybersecurity vulnerabilities that arise due to communication with neighboring vehicles or infrastructure [33, 34].

At the same time, there is no free lunch as far as connectivity is concerned. As argued in [19], forgoing connectivity comes at the cost of either more crashes or reduced throughput . In order to reap the benefits of connectivity, we need to tackle the challenges associated with it. Since more than 50% of all injury/fatal crashes occur at or near intersections [35], a first step could be to focus on making I2V technology at intersections economically viable. A fully instrumented intersection can cost around $25,000 [22]. Due to limited budgets and several pressing concerns, municipal bodies cannot go about instrumenting all intersections in their jurisdiction. However, crashes are not uniformly spread across all intersections. As observed in our study of intersections in Torrance, California [18], a few intersections account for most of the crashes. Thus, instrumenting even a few selected intersections can lead to substantial safety gains. Moreover, public-private partnerships are another increasingly popular option for funding such developments. A recent example is the partnership between Cavnue, a smart infrastructure company, and the state of Michigan to develop a 40 mile connected and autonomous vehicle corridor [36]. In addition, more research and investment is required to mitigate the cybersecurity challenges associated with this technology.

6 Conclusion

Autonomous vehicles (AVs) have the potential to lead us towards a safe transportation future. They can save thousands of lives that are lost each year due to impaired or rule-violating drivers. At the same time, faults in AV technology can claim lives, as evidenced by fatal crashes in the last few years. An accurate assessment of current AV safety capabilities is crucial to unlock the future safety benefits of this promising technology while mitigating its potential risks.

Due to the overwhelming diversity of driving contexts and rare nature of crashes, solely focusing on AV on-road testing data provides limited insights about its safety capabilities. In this paper, we discussed how human crashes that occur over billions of miles of human driving across diverse contexts can enhance our understanding of the benefits as well as limitations of current AV technology. We demonstrated that a context-aware analysis of human crashes not only allows us to compare the safety performance of AVs with human drivers but also sheds light on scenarios that are likely to remain challenging for them despite improvements in technology. Moreover, it enables us to translate the on-road testing performance of AV technology in specific ODDs to other jurisdictions where AVs haven't yet been tested. Finally, we discussed how such an analysis can inform policies and regulations geared towards the safe deployment of AV technology. Thus, we argue that the lessons learnt from human crashes can be invaluable in enabling AVs to bring about a safe transportation future.

Acknowledgement

This research was supported by National Science Foundation EAGER award 1839843, Berkeley Deep Drive, and California Department of Transportation.

References

[1] National Center for Statistics and Analysis. 2018 fatal motor vehicle crashes: Overview. (Traffic Safety Facts Research Note. Report No. DOT HS 812 826). Technical report, United States National Highway Traffic Safety Administration, 2019.

[2] Lawrence Blincoe, Ted R Miller, Eduard Zaloshnja, and Bruce A Lawrence. The economic and social impact of motor vehicle crashes, 2010 (Revised) (Report No. DOT HS 812 013). Technical report, United States National Highway Traffic Safety Administration, 2015.

[3] Santokh Singh. Critical reasons for crashes investigated in the National Motor Vehicle Crash Causation Survey (Traffic Safety Facts Crash Stats. Report No. DOT HS 812 115). Technical report, United States National Highway Traffic Safety Administration, 2015.

[4] Cameron F. Kerry and Jack Karsten. Gauging investments in self-driving cars. `https://www.brookings.edu/research/gauging-investment-in-self-driving-cars`, 2017.

[5] John M. Scanlon, Kristofer D. Kusano, Tom Daniel, Christopher Alderson, Alexander Ogle, and Trent Victor. Waymo simulated driving behavior in reconstructed fatal crashes within an autonomous vehicle operating domain. 2020. `https://storage.googleapis.com/waymo-uploads/files/documents/Waymo-Simulated-Driving-Behavior-in-Reconstructed-Collisions.pdf`.

[6] Andrew Hawkins. Serious safety lapses led to Uber's fatal self-driving crash, new documents suggest. *The Verge*, 2019.

[7] Tom Krisher and Olga R. Rodriguez. NTSB releases details in 2 crashes involving Tesla Autopilot. *Associated Press*, 2020.

[8] Phillip Koopman. A Reality Check on the 94 Percent Human Error Statistic for Self-Driving Cars. *Safe Autonomy Blog*, 2018. `http://safeautonomy.blogspot.com/2018/06/a-reality-check-on-94-percent-human.html`.

[9] L Lubbe, Hanna Jeppsson, Arian Ranjbar, Jonas Fredriksson, Jonas Bärgman, and Martin Östling. Predicted road traffic fatalities in Germany: The potential and limitations of vehicle safety technologies from passive safety to highly automated driving. In *Proceedings of IRCOBI conference. Athena, Greece*, 2018.

[10] Casualty Actuarial Society Automated Vehicles Task Force. Restating the National Highway Transportation Safety Administration's National Motor Vehicle Crash Causation Survey for Automated Vehicles. *Casualty Actuarial Society E-Forum, Fall 2014-Volume 1*, 2014.

[11] Elizabeth D Swanson, Frank Foderaro, Mikio Yanagisawa, Wassim G Najm, Philip Azeredo, et al. Statistics of Light-Vehicle Pre-Crash Scenarios Based on 2011–2015 National Crash Data. Technical report, Washington, DC: National Highway Traffic Safety, 2019.

[12] Waymo. Waymo Safety Report. 2020. `https://storage.googleapis.com/sdc-prod/v1/safety-report/2020-09-waymo-safety-report.pdf`.

[13] National Highway Traffic Safety Administration. National Motor Vehicle Crash Causation Survey: Report to Congress. *Technical Report DOT HS*, 811:059, 2008.

[14] Hannah Boland. Waymo's self-driving cars 'struggle to turn left and don't understand basic road features'. *The Telegraph*, 2018.

[15] Amir Efrati. Waymo's Foes: Left Turns and the Mean Streets of Phoenix. *The Telegraph*, 2017.

[16] Andrew Hawkins. Waymo's driverless cars hit a new milestone: 10 million miles on public roads. *The Verge*, 2018.

[17] Collision between vehicle controlled by developmental automated driving system and pedestrian. *National Transportation Safety Board, Washington, DC, USA, Tech. Rep. HAR-19-03*, 2019.

[18] Akhil Shetty, Hamidreza Tavafoghi, Alex Kurzhanskiy, Kameshwar Poolla, and Pravin Varaiya. Risk assessment of autonomous vehicles across diverse driving contexts. 2021. `https://doi.org/10.13140/RG.2.2.33166.28485`.

[19] Akhil Shetty, Mengqiao Yu, Alex Kurzhanskiy, Offer Grembek, Hamidreza Tavafoghi, and Pravin Varaiya. Safety challenges for autonomous vehicles in the absence of connectivity. *arXiv preprint arXiv:2006.03987*, 2020.

[20] Mike Isaac. Uber Suspends Tests of Self-Driving Vehicles After Arizona Crash. *The New York Times*, 2017.

[21] Ann Hsu, Farokh Eskafi, Sonia Sachs, and Pravin Varaiya. Protocol design for an automated highway system. *Discrete Even Dynamic Systems: Theory and Applications*, 2:183–206, 1993.

[22] Offer Grembek, Alex Kurzhanskiy, Aditya Medury, Pravin Varaiya, and Mengqiao Yu. Making intersections safer with I2V communication. *Transportation Research Part C: Emerging Technologies*, 102:396–410, 2019.

[23] James A Misener. Cooperative intersection collision avoidance system (CICAS): Signalized left turn assist and traffic signal adaptation. Technical report, 2010.

[24] Shai Shalev-Shwartz, Shaked Shammah, and Amnon Shashua. On a formal model of safe and scalable self-driving cars. *arXiv preprint arXiv:1708.06374*, 2017.

[25] David Nistér, Hon-Leung Lee, Julia Ng, and Yizhou Wang. The safety force field. *NVIDIA White Paper*, 2019.

[26] Andrea Censi, Konstantin Slutsky, Tichakorn Wongpiromsarn, Dmitry Yershov, Scott Pendleton, James Fu, and Emilio Frazzoli. Liability, ethics, and culture-aware behavior specification using rulebooks. In *2019 International Conference on Robotics and Automation (ICRA)*, pages 8536–8542. IEEE, 2019.

[27] Nidhi Kalra and Susan M Paddock. Driving to safety: How many miles of driving would it take to demonstrate autonomous vehicle reliability? *Transportation Research Part A: Policy and Practice*, 94:182–193, 2016.

[28] Nidhi Kalra and David G Groves. *The enemy of good: Estimating the cost of waiting for nearly perfect automated vehicles*. RAND Corporation, 2017.

[29] Andrew Hawkins. Everyone hates California's self-driving car reports. *The Verge*, 2020.

[30] Kyle Vogt. The Disengagement Myth, 2017. `https://medium.com/cruise/the-disengagement-myth-1b5cbdf8e239`.

[31] Philip Koopman and Beth Osyk. Safety argument considerations for public road testing of autonomous vehicles. *SAE International Journal of Advances and Current Practices in Mobility*, 1(2019-01-0123):512–523, 2019.

[32] Andrew Hawkins. Motional is now testing fully autonomous vehicles in Las Vegas. *The Verge*, 2021.

[33] Jin Cui, Lin Shen Liew, Giedre Sabaliauskaite, and Fengjun Zhou. A review on safety failures, security attacks, and available countermeasures for autonomous vehicles. *Ad Hoc Networks*, 90:101823, 2019.

[34] Ben Mccluskey. Connected cars–the security challenge. *Engineering & Technology*, 12(2):54–57, 2017.

[35] Intersection Safety. *Federal Highway Administration*. `https://highways.dot.gov/research/research-programs/safety/intersection-safety`.

[36] Michelle Grinnell. Michigan, Cavnue Creating Road of Future Between Ann Arbor and Detroit. *Michigan Economic Development Corporation*, 2020. `https://www.michiganbusiness.org/press-releases/2020/08/michigan-cavnue-creating-road-of-future-between-ann-arbor-and-detroit/`.