Journal of Safety Research 63 (2017) 57–60

Contents lists available at ScienceDirect

Journal of Safety Research

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

Rage against the machine? Google's self-driving cars versus human drivers

Eric R. Teoh, * David G. Kidd

Insurance Institute for Highway Safety, United States

ARTICLE INFO

Article history:
Received 9 June 2017
Received in revised form 31 July 2017
Accepted 17 August 2017
Available online 26 August 2017

Keywords:
Autonomous vehicle
Self-driving
Driving automation
Motor vehicle crashes
Highway safety

ABSTRACT

Introduction: Automated driving represents both challenges and opportunities in highway safety. Google has been developing self-driving cars and testing them under employee supervision on public roads since 2009. These vehicles have been involved in several crashes, and it is of interest how this testing program compares to human drivers in terms of safety. *Methods:* Google car crashes were coded by type and severity based on narratives released by Google. Crash rates per million vehicle miles traveled (VMT) were computed for crashes deemed severe enough to be reportable to police. These were compared with police-reported crash rates for human drivers. Crash types also were compared. *Results:* Google cars had a much lower rate of police-reportable crashes per million VMT than human drivers in Mountain View, Calif., during 2009–2015 (2.19 vs 6.06), but the difference was not statistically significant. The most common type of collision involving Google cars was when they got rear-ended by another (human-driven) vehicle. Google cars shared responsibility for only one crash. *Conclusions:* These results suggest Google self-driving cars, while a test program, are safer than conventional human-driven passenger vehicles; however, currently there is insufficient information to fully examine the extent to which disengagements affected these results. *Practical application:* Results suggest that highly-automated vehicles can perform more safely than human drivers in certain conditions, but will continue to be involved in crashes with conventionally-driven vehicles.

1. Introduction

Since the first automobile, many vehicle functions have been automated, such as starting the engine, shifting transmission gears, and opening doors and windows, but the actual driving task, fundamentally, has changed very little. Automating the driving task represents a huge safety potential, but it also presents many challenges. Computers don't drive while impaired by alcohol and can process larger amounts of information to make quicker decisions compared with humans. On the other hand, development of these technologies takes a long time since the driving environment is complex, and the human and economic costs of errors are high. Another major challenge, particularly of early systems, is how driving responsibility is shared and exchanged between human and machine. For instance, current systems that automate part of the driving task, such as Tesla's "Autopilot," do not assume responsibility for monitoring the driving environment. The human drivers are assumed to remain attentive and to be able to take over vehicle control at any moment should the automated systems fail in any manner. Even with systems that do assume some responsibility, humans may over-rely on such systems, may fail to adequately take over vehicle control, or have other difficulties using automation as observed in other domains (Billings, 1991). One way to solve these issues is to automate all parts of the driving task and remove all possible direct human control. This is the approach favored by Google's self-driving car project, and is what the SAE International would classify as SAE level 4 or 5 automated driving systems (SAE, 2016).

Google's self-driving car project, now known as Waymo, has been performing supervised autonomous driving, in which the vehicle systems control all aspects of the driving task with Google engineers supervising and re-taking vehicle control if necessary, on public roads since 2009, beginning with a fleet of modified Toyota Prius cars. In 2012, Google switched to a fleet of modified Lexus RX450h SUVs and shifted testing from freeways to more urban roads, mostly in Mountain View, Calif. (Waymo.com). Google began testing a low-speed prototype vehicle on public roads in 2015. Google also has been testing vehicles on public roads in Austin, Texas; Kirkland, Washington; and metro Phoenix, Arizona. Alphabet Inc., the parent company of Google, spun off the self-driving car project into an Alphabet-umbrella company called Waymo in 2016, but we refer to this program and its vehicles colloquially as the Google car throughout this paper. A natural part of testing on public roads is involvement in crashes, and Google cars have been involved in several. Google has reported crash events, regardless of severity or fault, in a series of monthly activity reports (Google, 2015-2016) and, in California, to the California Department of Motor Vehicles (DMV).

* Corresponding author.
E-mail address: eteoh@iihs.org (E.R. Teoh).

http://dx.doi.org/10.1016/j.jsr.2017.08.008
0022-4375/© 2017 Published by Elsevier Ltd.

Two studies have investigated the crash experience of Google cars and compared them with conventionally-driven vehicles. Schoettle and Sivak (2015) studied crash rates reported by self-driving car test programs from Google, Delphi, and Audi, but all the crash events and the vast majority of mileage were from Google cars, with Delphi and Audi contributing less than 5000 miles compared with Google's 1.2 million miles, so this was effectively a study of the Google car experience. They concluded that self-driving vehicles have crash rates more than double those of conventional vehicles, albeit not a statistically significant difference. They arrived at this conclusion by inflating rates of police-reported crashes in the United States to account for underreporting to police, using estimates of the magnitude of underreporting from other studies and comparing them with self-driving car crash incidents (most of which were very minor and not police-reportable). This form of comparison, however, assumes that the distribution of unreported crash types of conventional vehicles is the same as the distribution of low-severity crashes of self-driving cars. The study investigated differences in crash types, but the comparison methods were imprecise. For instance, the authors compared self-driving cars in terms of the proportion of crashes that were front-to-rear. However, this comparison is flawed because approximately half of human-driven vehicles involved in front-to-rear crashes were struck from behind, whereas all the Google cars involved in such crashes were struck from behind.

Blanco et al. (2016) found Google cars generally have lower crash rates than conventional vehicles by analyzing data from the Second Strategic Highway Research Program (SHRP2) naturalistic driving study. The SHRP2 study involved over 3000 participating drivers across six study sites throughout the United States during 2010–2013. Each participant drove a vehicle outfitted with sensors and video to record all driving for up to 3 years. The advantage of using naturalistic data is that all events where the subject vehicle strikes another object are recorded, even those that are less severe. Crash events in the SHRP2 data were classified as level 1 (potentially serious crash and likely police-reported) to level 4 (e.g., a tire bumping a curb and almost certainly not police-reported). Blanco et al. mapped the Google car crash events onto the four-level classification system and then compared rates within level. These should be fair comparisons in terms of minimizing reporting differences, but geographical differences exist. To address the geographic variation issue, Blanco et al. looked at police-reported crash rates in some of the SHRP2 counties and in similar counties where Google cars primarily operate. Since this brings back the issue of underreporting to police, they used three estimates of underreporting to inflate the police-reported crash rates. Also, they restricted the Google car crash events used in calculation of rates to those of level 1 or 2, as those are more likely to be reported to police. However, virtually all the Google car crashes were in Mountain View, suggesting that two counties form an unnecessarily broad comparison group.

The purpose of this study was to examine differences in the crash experience of Google self-driving car testing and that of conventional passenger vehicles driven by humans using straightforward methods to address reporting and geographical differences.

2. Methods

Two general strategies for comparing the crash experience of Google cars during testing and human passenger vehicle drivers were used in this study: comparing crash rates and comparing patterns of crashes. The primary difference between this analysis and the earlier studies is that no assumption was made about the distribution of unreported crash incidents when comparing rates. Rather than inflating reported crash rates for the comparison group (human drivers), the crash rates for Google cars were restricted to those deemed to be police reportable.

Data were obtained from a variety of sources, as outlined in Table 1. California requires every crash involving an automated vehicle that results in any property damage, injury, or death to be reported to California DMV. This reporting was standardized with the creation of the OL316 form in 2014. Google has been providing written narratives of each crash (Google, 2015–2016) in all states where testing occurs and has reported crash incidents that did not result in any damage, suggesting its reporting of such incidents is complete. Crash incidents and vehicle miles traveled for Google cars were counted only while the vehicles were operated in autonomous mode. For each crash incident, narratives were used by the authors to code type of collision, injury severity, and whether the crash likely met the California threshold for police crash reporting (fatality, injury, or $750+ in property damage) during the study period. It was unknown if police crash reports were filed for Google car crashes.

Table 1
Data sources.

Data elements	Sources
Google car crash incidents	Google monthly activity reports; California OL316 forms required for each self-driving vehicle crash occurring after May 2014
Google car vehicle miles traveled (VMT) in autonomous mode in all states, combined	Google monthly activity reports
Police-reported crashes in California	2010–15 California Statewide Integrated Traffic Records System (SWITRS)
Police-reported crashes in the United States	2010–15 National Automotive Sampling System General Estimates System (NASS-GES)
VMT in California	2009–14 California Department of Transportation Highway Performance Monitoring System annual reports
VMT in United States	2010–15 Federal Highway Administration Highway Statistics Series (VM-1 form)
Naturalistic driving crash incidents	SHRP2 naturalistic coded/video data version 3.3 accessible via InSight website interface

California vehicle miles traveled (VMT) data were provided as average miles per day, so these were extrapolated to annual totals by multiplying by 365 (or 366 for the leap year, 2012). VMT data and California police-reported crash data were available at the city/county level. At the time of this study, California VMT data were not available for 2015 so the 2014 figures were used as estimates in calculating crash rates for 2010–2015 in aggregate. California VMT included all vehicle types, whereas Google's reported VMT was specific to the test vehicles. National VMT estimates were available for passenger vehicles, but the total figure was taken for consistency with California. Passenger vehicles comprise the vast majority of VMT.

Based on numbers of vehicles by city stated in monthly reports, Google began expanding beyond California largely in 2016, so the comparisons with California crash data were restricted to pre-2016. Moreover, all of Google's crash incidents before 2016 occurred in the Mountain View area. Rates of Google car crashes judged likely to be reportable to police were compared to rates of passenger vehicles in crashes reported to police in Mountain View, Santa Clara County, California, and the United States. Crash rates were calculated per million vehicle miles traveled.

Crash types were compared by tabulating observed counts of police-reportable (and all) crashes for Google cars and expected counts had they crashed with the same distribution as police-reported crashes of human-driven passenger vehicles in Mountain View. The latter, for each crash type, was calculated by multiplying the rate of that type of crash per VMT in Mountain View by the VMT reported by Google. This allows for inspection of the distribution of crash types of human drivers and provides for comparisons of Google car crash rate for specific types of crashes. Crash types also were compared using an analysis of SHRP2 data. Crash rates for drivers in SHRP2 and for Google cars were tabulated by type of crash. Crashes in SHRP2 are defined by contact and vehicle