Automated Car Woes—Whoa There!

by Peter G. Neumann

Editor's Introduction

With all the growing interest in automated cars and driverless cars and recent accidents involving them, we thought we would turn to Risks founder Peter G. Neumann for perspective. Neumann has been moderating the ACM Risks Forum (risks.org) since 1985, and has accumulated a vast trove of experience in the ways that automated systems can not only provoke but enable mishaps. Here he offers an inventory of how automotive automation systems can fail and concludes that the greatest threats to safety come from human tendencies such as being distracted at the moment the system needs an intervention, or having their skills as operators grow rusty over time because the automation is pretty good most of the time.

—Peter J. Denning, Editor in Chief

Automated Car Woes—Whoa There!

by Peter G. Neumann

If you are feeling queasy about the prospect of driverless cars, heavily automated cars, and automated highways, you have good reasons to be. New risks invariably accompany new technology. Let's take a look at where the risks are, and what it might take to reduce them.

Once upon a time, horses were the primary means of transportation. Crude roads and narrow paths accommodated equestrian and foot traffic. When horse-drawn carriages became popular, roads were widened and improved to be navigable in winter and during storms. Then in the 1880s came the first automobiles. They moved much faster and required better roads. They were initially called horseless carriages. The term automobile started coming into use around 1900. Automobiles were themselves an automation of a previous technology. Next came Henry Ford, who mass-produced Model T cars for business and pleasure from 1908 to 1927, filling the roads with millions of vehicles.

Since then automotive companies have been on a relentless quest to make cars easier and safer to drive, and government transportation authorities have worked to design roads and regulations that improve safety. In the 1960s, Ralph Nader made automotive safety a political imperative, with the publication of his book, *Unsafe at Any Speed*. This 1965 best-seller attacked the safety record of American automobile manufacturers in general, and the first-generation of the Chevrolet Corvair in particular. His trenchant critique sent shockwaves across the U.S., with ripple effects in Europe and much of rest of the world.

Beginning in 1968 with Volkswagen, automobiles began incorporating embedded systems that automatically controlled part of the cars, improving safety and ease of driving. Examples include anti-lock brakes, navigation systems, drive by wire, adaptive cruise control, airbag control, and automatic maintenance. Among the more recent innovations has been a control system that will parallel park a car. These systems today rely extensively on embedded computers, and the entire vehicle has become controlled via a local network connecting essentially everything. This implied that the microcomputers had to work seamlessly together, and made computerized diagnostics a centerpiece of the role of auto mechanics.

Automobile makers have been also experimenting with voice-controlled interfaces on the (relatively untested) theory that hands-free operation might be less distracting. These interfaces work fairly well when a single person is giving the commands in a normal voice. However, automobile hazards can come up rapidly. The driver's voice can suddenly become strained and compete with passengers' yells. Also, a new driver may not be recognized unless having been the person who trained the software. There's quite a lot of artificial intelligence (AI) in that part of the human interface, and significant potential for the computer recognizing some threats before the would-be driver could—but there is also significant

potential for it reacting inappropriately. How the voice controlled system will react in those cases is unpredictable.

When I think another driver is about to put me in great danger, I blow my horn and shout a few expletives—and my passengers may be doing likewise. This could be a likely time for the system to say, "Sorry, I cannot understand that command." There is something to be said for the days-gone-by when the horseman could avoid an accident by saying, "Whoa there!" to the horse.

Now we stand at the threshold of a new age, with the somewhat overhyped prognostication of the advent of driverless cars and more advanced automated computer-controlled driver assists. Google and Tesla have been among the leaders in developing prototypes and testing them in live traffic. Google aims at self-driving cars, even though a surrogate driver is currently present for testing purposes. Tesla focuses on semi-autonomy, with computerized assists for drivers rather than completely automated vehicles. Other car makers are exploring various points in the spectrum heading toward extensive autonomy. Preliminary data suggest that these cars have the potential to be much less accident prone than standard cars operated by human drivers. According to their promoters, driverless cars combined with "smart highways" could dramatically reduce accident rates and traffic congestion, which is a very attractive prospect.

In the wake of a recent mishap with a Tesla S car in Florida, which killed its passenger, critics have begun to question whether AI software in the car controls might be too risky. Ironically, Elon Musk, the founder of Tesla Motors, took a stand recently that AI may be dangerous for humanity. Is his prediction showing up in his own cars? Are passengers entitled to trust that these cars will be safe when under automatic control?

(Perhaps what might provide a clear demonstration of the safety of the emerging automated vehicles would be a "Whoa's Ark," in which two of each automated vehicle model are pitted against one another in a full-scale, bumper-car-like competition free-for-all.)

In addition to those who worry the software contains bugs and design flaws that endanger people on the road, there are those who worry Internet-connected cars can become susceptible to malware, hackers, and intruders. Very little local cybersecurity technology has been incorporated into vehicles at this time. Furthermore, almost all the components are now becoming electronically controlled on a common network—without real concerns for either internal or external security. Thus, it is not surprising in this decade we have seen numerous demonstrations of vulnerabilities that could enable malevolent individuals to take over the controls on modern cars remotely—for example, through the wireless maintenance port, without any need to use the Internet!

The Tesla S incident demonstrates the risk of people trusting software that has yet to establish a track record of safety. The Tesla driver in Florida died because he gave too much credence to the Tesla S Autopilot. (He did so quite publicly. A selfie on the Web shows him smiling in the driver's seat with no hands on the wheel.) In a separate incident, another Autopilot-trusting driver was injured when his Tesla X bounced off the siding and flipped over while on autopilot. (The term "autopilot" is clearly a misnomer, as Tesla insists the driver must necessarily remain in the loop. Indeed, *Consumer Reports* and others have urged Tesla to eschew the term, and to *require* the surrogate driver to keep hands on wheel at all times.) On the other hand, Google's cars are intended to be self-driving, once they are deemed

safe enough. Even so, Google cars have had a few accidents when they were rear-ended because their very conservative programming was a mismatch for the tailgating driver behind, who perhaps expected the car in front of them to run a red light in the tradition of New York city taxi drivers.

John Quain notes there is significant evidence that a driver behind the wheel may not be ready to take over from the autopilot quickly enough to avert a disaster: "Experiments conducted last year by Virginia Tech researchers and supported by the national safety administration found that it took drivers of Level 3 cars—in which the driver can fully cede control of all safety-critical functions in certain conditions—an average of 17 seconds to respond to takeover requests. In that period, a vehicle going 65 mph would have traveled 1,621 feet— more than five football fields [1]."

Quain also notes Google discovered (on video-cam) that its surrogate drivers-in-waiting were often far too distracted to react quickly enough. He added, "more conventional automakers have designed their systems to take control of the car for only a few seconds at a time; the driver must be ready to resume command at any time."

The media seem to place a lot of faith in the role of AI to make driverless cars safer. Their conception is all the data collected by cars about fine details of road conditions and traffic patterns could be exploited to keep the cars away from hazards that humans might not even see. This is a hypothesis that will take years of experience to decide. To date, AI systems have performed well in the situations for which they were trained, but demonstrated brittleness or downright stupidity in other situations. Many vehicle mishaps are triggered by surprises that were not part of the training scenarios. It is too early to tell whether automated systems can learn enough to deal with surprises better than humans.

Even assuming we take a cautious view about the potential for AI in vehicles, it appears the *control* software in today's cars is mostly *not* AI. (This may be somewhat simplistic, as the voice software is typically not considered part of the control software—although it can trigger deleterious effects.) Typical automotive software combines geolocation, sensor and video data, and carefully engineered algorithms to control an embedded system. The software control mishaps to date seem not directly attributable AI.

There is an accelerating practice in the software industry to push out upgrades without user intervention; this is often done in the name of security to quickly remediate vulnerabilities, but could easily compromise automotive safety. It is undoubtedly more sensible to install updates when the car is shut off (not driving), but many users will want to defer updates until they are satisfied with the level of testing. Moreover, even if the owner is ready to trust the safety of an update, it may be desirable to defer the update until a minimally disruptive moment. There is also a risk the upgrade itself is botched because it creates new vulnerabilities, is co-opted by an attacker, or is garbled in transmission. Many upgrades actually require remedial assistance.

The bottom line here must be one of extreme caution. There is much research yet to be done if automated cars can be deemed safe. Tesla insists their newer cars are still experimental and require constant awareness on the part of the responsible person in the vehicle. Google makes enormous efforts to test new software extensively and to monitor all activities of their cars and surrogate drivers. Despite the recent accidents, the safety record of these cars seems to be vastly better than that of conventional cars under the control of distracted, aggressive, or poor drivers.

Driver education is about to take a big leap in spite of claims that less education will be needed. Drivers may not need education on risks that have been eliminated by automation, but they need to be educated in how to respond to the new risks that invariably accompany a new technology. In addition, testing, evaluation, and regulation are going to be crucial. Evaluation will be particularly tricky, because of many varying conditions, assumptions, and claims of developers. Even worse, some developers have a history of hiding some of their internal analyses that assign somewhat arbitrary low costs to critical factors. Sooner or later flaws known internally are found out and turn out badly—consider the Ford Pinto gasoline-tank risks and the Volkswagen emissions fiasco.

I have been moderating the ACM Risks Forum (risks.org) for 31 years. We collect details about problems, mishaps, and catastrophes caused by computer hardware, software, or unexpected human behavior. Some mishaps were fatal, caused injury, led to costly damage, ruined reputations, or were simply annoying. One of our big lessons is that these failures have many causes—for example, incomplete requirements specifications, flawed system design, sloppy implementation, poor human interfaces, human error, malicious interventions, and more. Any system containing software controls should be developed conservatively, with extensive testing and evaluation. Regulation is also needed. The process must be iterative and its managers ready to correct any problems discovered. Even so, the process is not perfect and there will always be a residual risk of failure. No vehicle will ever be risk free. And we can be confident that the completely automated highway will be enormously complex, which will compound all of the problems noted here and pile on the requirement for coordination among all neighboring vehicles and the entire operational environment.

Tesla, Google, and all the others seeking to compete in this marketplace need to pay much greater attention to the history of research and development in autonomous systems. In response to a draft version of this article, Don Norman noted there is a large body of relevant literature relating to human attention characteristics, vigilance, and situational awareness—which has not been read by today's developers. He believes partial automation is inherently dangerous because it exacerbates unsafe human tendencies [2,3]. He prefers full automation: "To think otherwise is to ignore decades of solid research from the psychology and human factors fields and the National Academy's Human Systems Integration board. And there is no way to overcome it. The better the automation, the more dangerous it becomes. It has to be full automation, not this silly National Highway Safety Administration level 3."

Don offers this wisdom: The most difficult part of automation is not the technical part, it is dealing with people, novel situations, and unexpected events. He adds, "We know two things about unexpected events: One, they will happen. Two, when they do happen, they will be unexpected."

In addition, I have long maintained the subordinate goal of eliminating system administrators in the context of highly autonomous systems entails some very perverse risks. When something does go badly wrong, there are likely to be very few system administrators with sufficient training and experience to diagnose and repair what has collapsed. And here we would be expecting the mechanics to become supremely skilled system administrators!

Overall, let everyone beware—whoever might believe in the infallibility of an automated automotive system, including drivers, developers, marketers, and mechanics. Even if the development efforts are done wisely and cognizant of the reality that these are total-system problems that encompass research, development, engineering, human interfaces, operational issues, and much more, totally self-driving

vehicles and the automated highway will still be a long time coming. If they are not done wisely, it will be an even longer time coming.

About the Author

[Peter G. Neumann](#), Ph.D., is a senior principal scientist at SRI International Computer Science Lab, moderator of the ACM Risks Digest ([risks.org](#)), and editor of the *Communications of the ACM* “Inside Risks” series, for which Neumann has an article on the broader risks of automation in the October 2016 issue.

References

[1] John B. Quain. The autonomous car vs. human nature, a driver behind the wheel may not be ready to take it. *The New York Times*. July 8, 2016

[2] S. M. Casner, E. L. Hutchinson, and D. Norman. The challenges of partially automated driving. *Communications of the ACM* 59, 5 (May 2016).

[3] Donald A. Norman. The human side of automation. February 24, 2015. Keynote address for Automated Vehicles Symposium, 2014. Published in *Road Vehicle Automation 2* (Springer, 2015.)

Note

Relevant [Ubiquity blogs](#) and [articles](#) addressing this topic include:

- [“What About an Unintelligent Singularity?”](#) Peter J. Denning,
- [“A Shortage of Technicians”](#) Peter J. Denning
- [“Your Grandfather’s Oldsmobile—NOT!”](#) Ted Lewis
- [“Can Robots Be Trusted?”](#) Lewis Perelman
- [“The Future of Technology and Jobs”](#) Arun Mumar Tripathi

DOI: 10.1145/2974062