What Do We Tell the Drivers? Toward Minimum Driver Training Standards for Partially Automated Cars

Stephen M. Casner, National Aeronautics and Space Administration, USA,
and **Edwin L. Hutchins**, University of California San Diego, USA

Each year, millions of automobile crashes occur when drivers fail to notice and respond to conflicts with other vehicles, bicyclists, and pedestrians. Today, manufacturers race to deploy automation technologies to help eliminate these mishaps. To date, little effort has been made to educate drivers about how these systems work or how they affect driver behavior. Driver education for automated systems amounts to additional pages in an owner's manual that is known to be a seldom-used glove box reference. In this article, we review the history of automation deployed in the airline cockpit decades ago. We describe how automation helped avoid many common crash scenarios but at the same time gave rise to new kinds of crashes. It was only following a concerted effort to educate pilots about the automation, about themselves, and about the concept of a human-automation team that we reached the near-zero crash rate we enjoy today. Drawing parallels between the automation systems, the available pilot and driver research, and operational experience in both airplanes and automobiles, we outline knowledge standards for drivers of partially automated cars and argue that the safe operation of these vehicles will be enhanced by drivers' incorporation of this knowledge in their everyday travels.

Keywords: advanced driver assistance systems, cockpit automation, standards, traffic safety, training

For the great majority of drivers, car automation technologies such as pre-collision warning, automatic emergency braking (AEB), and hands-free driving systems have remained the topic of news stories that report the most noteworthy advances and spectacular failures. Hands-on (or hands-off) experiences with these systems have been limited to early adopters who paid a premium to gain first access to these futuristic technologies. In 2019, that is set to change, as it is the year that advanced driver assistance systems (ADAS) begin their migration from optional equipment in top-of-the-line cars to standard equipment in lower-priced models (Consumer Reports, 2018). Hands-free highway pilots will begin to appear in more affordable cars offered by legacy car manufacturers. For a great many drivers, 2019 will mark a revolution in driving on a scale never before seen.

The question of how drivers will cope with this transition from strictly manual driving to a driving task that is shared with automation is of great interest to those who work in safety-related fields. To date, the industry's approach to preparing drivers for this transition is a passively aggressive one. At a minimum, drivers will be provided with additional pages in their driver manual that outline the button-pushing procedures required to operate these systems—along with a detailed list of the situations in which the systems may not function as anticipated. While these are necessary first steps, human factors researchers know that there is more to the story about what happens when humans work cooperatively with machines. What we have learned comes from the human factors research that began in the late 1970s when similar automation was introduced in the airline cockpit. That early research yielded a number of unexpected findings. We found that pilots were sometimes surprised by the behavior of the automation, unable to predict what a complex system would do next. Ironically, being able to predict what the automation would do next seemed to require more

---

Address correspondence to Stephen M. Casner, NASA Ames Research Center, National Aeronautics and Space Administration, Mail Stop 262-4, Moffett Field, CA 94035, casner@gmail.com.

Journal of Cognitive Engineering and Decision Making
201X, Volume XX, Number X, Month 2019, pp. 1–12
DOI: 10.1177/1555343419830901

knowledge about how the automation works than we originally anticipated. We saw how working with automation sometimes taxed pilots' ability to pay attention and manage distractions—ones often introduced by the automation itself. We saw how automation changed the job of flying in fundamental ways. Today, it is standard practice to provide pilots with a basic understanding of humans, machines, and what happens when the two are combined. All the while, we enjoy a historic low in the aviation crash rate.

Over the past 10 years, human factors researchers have begun to study these same phenomena with drivers in partially automated cars. It comes with little surprise that researchers are seeing many of the same human factors issues arise. Forty years the wiser, we have an opportunity to use what we have learned to help prepare drivers as they transition to the world of partially automated driving. Yet, little industry attention has been aimed in this direction.

In this article, we gather the lessons learned from 40 years of human-automation research in both planes and automobiles. We conclude with a first attempt at a set of standards for what drivers should know before they operate a partially automated car (whether purchased, rented, or borrowed) and consider how we might go about delivering these concepts to drivers who may or may not understand that they even need them.

AIRPLANES

In the early days of air travel, most airline crashes happened when flight crews either lost control of the airplane or unknowingly flew into terrain or another aircraft (Boeing, 2017). To help reduce these mishaps, the industry turned to automation technologies: onboard systems that promised to guard against them. In 1974, the Federal Aviation Administration (FAA) mandated the use of a system that alerts flight crews when they are headed toward rising terrain (Breen, 1999). In 1993, the FAA mandated a similar system that detects other aircraft that have come into dangerous proximity (FAA, 2011). By 1983, some airplanes were equipped with systems that would prevent pilots from placing the airplane in dangerous attitudes that could lead to loss of control. By the same year, autopilots became capable of semi-autonomously tracking large portions of the pre-planned route, freeing pilots from the tedium of scanning instruments and operating controls, allowing them more time to remain on the lookout for anything gone wrong.

The hope was that these new systems would simply eliminate the occasional mishaps from an otherwise safe process of flying an aircraft. However, the automation also changed the task of flying the airplane in surprising and fundamental ways. We began to see new problems emerging in the way pilots interacted with the automation along with new kinds of crashes that arose from these problems.

The automation, powered by millions of lines of computer code, sometimes did things that pilots didn't expect and sometimes it neglected to do things that they did expect: what the aviation industry came to refer to as *automation surprises* (Sarter, Woods, & Billings, 1997). "What's it doing now?" and "What's it going to do next?" became commonly asked questions in the cockpit. Misunderstandings about what the automation was configured to do next resulted in crashes and lives lost (Sarter & Woods, 1995).

With warning systems in place to call out dangerous circumstances, pilots sometimes used them not as a backup but rather a primary means of monitoring the progress of the flight. In 1988, a flight crew chatted about a non-flight-related topic just before takeoff. Little did they know, they had forgotten to set the airplane's wing flaps before they departed. The airplane stalled and crashed moments after takeoff, killing 2 crew members and 12 passengers. The airplane had a warning system onboard designed to automatically detect and call out mis-set flaps. The problem: the system wasn't working that day (National Transportation Safety Board [NTSB], 1989).

Pilots sometimes seemed unsure who was the higher and more reliable authority: the humans or the automation. In some cases, crews struggled to take back control of an airplane that was being automatically controlled (Ministry of Transport, 1996).

These unanticipated problems prompted the industry to revisit the idea that automation systems alone could be used to eliminate air disasters. How could engineers design around the possibility of a misunderstanding about how a complex system works? Or the human tendency to pay less attention to a job that is now being handled by someone (or something) else? Or the confusion surrounding the question of whether or not to trust an automation system in any given situation? There is no mechanical lockout device to protect pilots from their own thoughts, beliefs, and attitudes. Thus, the industry arrived to the next solution: to educate the pilots about the automation and the fundamental changes in the job of flying that it had introduced.

The airline industry took the standard first step toward educating the users of any product: they placed extra paragraphs in the operator's manual. The 1970 Boeing 727-200 Airplane Operations Manual contained 342 pages. A typical 1983 Boeing 757 manual contained 1,500 pages. Today, an Airbus A-320 Flight Crew Operations Manual contains roughly 2,700 pages. While the industry understood manuals to be a handy reference, the literature on instruction manual usage inspired little confidence that they would be carefully read and frequently used (Carroll & Rossen, 1987; Rettig, 1991). The daunting length of the manuals would also tax any human's ability to memorize and recall all of the material contained in them.

The industry next turned to research aimed at understanding how the introduction of automation systems was affecting pilots and the job of flying—and pilots were involved along every step of the way. Meetings and workshops were held that brought researchers and pilots together. Pilots began to advise researchers in their studies, and soon became co-investigators, sometimes earning advanced degrees to better enable them to do research. Researchers headed to nearby airports and learned to fly to help them better understand the problems they had been charged with studying. The end result of these collaborative efforts was to formulate and pass on three kinds of knowledge that the industry soon felt was indispensable to any pilot who stepped into this new kind of airline cockpit.

Knowledge About the Automation

Studies of flight crew uncovered basic ways in which pilots were confused by automation (Sarter & Woods, 1992, 1995). These findings were used to propose useful mental models of the automation that didn't require pilots to sift through millions of lines of computer code (Hutchins, 2007). Teaching methods were proposed (Feary et al., 1998) and books began to appear that helped airline pilots (Bulfer, 2018) and pilots still in training (Casner, 2013) to better understand the basics of how automation works. Empirical studies were conducted to measure what pilots took away from these training efforts (Casner, 2003, 2005).

Knowledge About Pilots

Studies that revealed how pilots' attention naturally drifted when automated systems were used prompted an industry-wide focus on the "out-of-the-loop phenomenon" (Endsley & Kiris, 1995; Wiener & Curry, 1980). The role of the pilot who was not handling the flight controls, previously named the "pilot not flying (PNF)" was redefined to a more active role and renamed the "pilot monitoring" and techniques were taught to help pilots be better monitors (Flight Safety Foundation, 2014). Human response to alerts and alarms were discussed widely, including the topics of surprise and startle (Thackray, 1965; Warrick, Kibler, & Topmiller, 1965), human reaction to false alarms (e.g., the "cry wolf" effect) (Breznitz, 1984), and takeover abilities under duress (Casner, Geven, & Williams, 2013).

Knowledge About the Changed Task of Flying

Researchers looked at ways to help pilots understand their new role as team players in a new cockpit consisting of both humans and automation (Kanki, Anca, & Helmreich, 2010; Sarter & Woods, 1997). Standards were developed to help prepare pilots to work in a team environment in which multiple agents had valuable inputs and in which any person or any thing could be right or wrong at any given time.

Pilots were reminded that they were in charge of and ultimately responsible for the outcome of the flight. That automation was a tool placed in their service, not the other way around. Pilots who were biased toward trusting the automation over their own judgment (Parasuraman & Riley, 1997) were given "turn-it-off training" during their initial exposure to automated aircraft.

Did these efforts pay off? Figure 1 shows the airline crash statistics from 1960 to the present. The precipitous drop in the crash rate through the

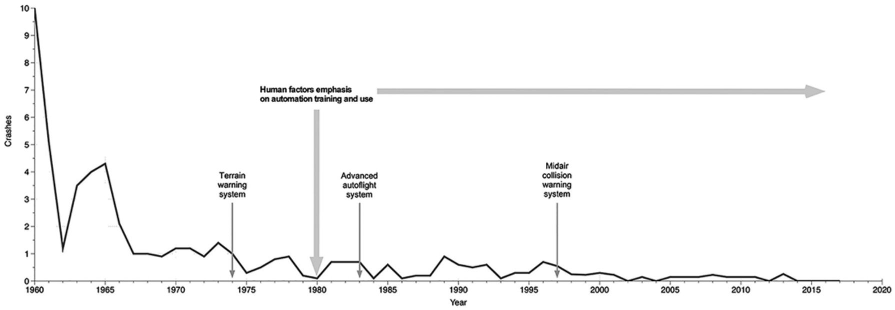

Figure 1. U.S. air carrier crashes by year.

1960s can be directly related to the introduction of radar and the jet engine to commercial aviation. These two innovations quickly reduced the number of mid-air collisions, flights into terrain, encounters with hazardous weather, and mechanical failures (Nolan, 2010). But in the late 1960s, progress in safety leveled off. The discrete events that were the introduction of the two warning systems did not produce immediate drops in the crash rate. Even as the prevalence of advanced autopilot systems, introduced in 1983, reached roughly 50% of the airline fleet by 1996, the crash rate remained the same over that 13-year period. What coincided with the eventual flattening of the Figure 1 graph in the late 1990s and early 2000s were the human factors efforts we described above. Although we saw few further changes to the automated cockpit, no regulatory changes, and no abrupt changes in the pool of available pilot candidates, airlines incorporated human-automation research findings into training programs such as crew resource management (CRM) and threat and error management (THEM). The concept of an effective but sometimes fallible human-automation team became widespread throughout the industry. Today, the airline industry still has problems related to pilot-automation interaction and crashes occasionally still happen as a result of them (NTSB, 2014). But through industry efforts to design automated systems, and then to provide pilots with a more complete understanding of their likely effects, these crashes have become rare events.

CARS

Today, we are seeing history repeat itself in cars. As a response to the many crashes that happen as a result of unintended lane and road departures, lane-keeping assist systems have been installed. Intersection, rear-end, and pedestrian collisions are being addressed with AEB systems and pre-collision warning systems. Highway pilot systems allow drivers to take their hands and feet off of the controls and allow the automation to assume continuous control. Highway pilot systems come with the hope that the overall behavior of a vehicle controlled by automation will be more stable and predictable (Watzenig & Horn, 2017).

But as manufacturers race to get these systems installed in cars, studies are already documenting how drivers allow their attention to drift (Barry, 2018; Carsten, Lai, Barnard, Jamson, & Merat, 2012; Hergeth, Lorenz, Vilimek, & Krems, 2016; Ledesma, Montes, Poó, & López-Ramón, 2015; Miller & Boyle, 2017), over-rely on or ignore alerts (Naujoks, Kiesel, & Neukum, 2016), don’t understand how the automation works (McDonald, Carney, & McGhee, 2018), and often make dangerous assumptions about what the automation is capable of doing and about what the driver is now free to do once the automation is turned on (Stewart, 2018). From the perspective of the aviation industry, it’s the 1980s all over again.

It is reassuring to see that meetings between driving researchers, car companies, insurance agencies, government agencies, and law firms are taking place and that these issues are being discussed. But the final step is missing. To date, there are few industry-wide plans in place to provide drivers with the sort of knowledge and awareness of the likely effects of automation that was provided to pilots. We understand that automobile drivers are more loosely selected and cursorily trained than pilots but none of the differences between pilots and drivers suggest that drivers should not receive training on the automation in their cars. Comparing the flying studies done years ago with the driving studies being done today, we see many similarities between these two groups when automation is introduced in either vehicle. We have already seen automation-related crashes happen in laboratory simulators and on the roads. And while survey studies are telling us that drivers across all age groups *want* to learn more about their new cars (Loeb, Belwadi, Shaikh, Shen, & Ward McIntosh, 2018), driver performance studies are telling us that teaching drivers about their cars has beneficial effects. Studies by Barg-Walkow and Rogers (2016) and Körber, Baseler, and Bengler (2018) found that even the simplest of introductory advice given to drivers had lasting effects on the way they perceived and used the automation.

In the remainder of this paper, we gather the lessons learned in the airline industry to outline a set of knowledge standards for drivers of partially automated cars. Here is what we believe every driver should know about their partially automated car and why we will be doing them a disservice if we don't tell them.

Knowledge About the Automation

Until now, the need for drivers to know about how cars work has been minimal. The basic controls of the car can be quickly mastered and there is little need to understanding the inner workings of any car component. We argue that car automation is changing that. Furthermore, advances in the capabilities of these systems will come rapidly and keeping up with them will be difficult for drivers. More challenging still is the situation in which drivers rent vehicles that contain this equipment or are tossed the keys to a neighbor's car—leaving them little time to inform themselves about how the car works.

Functions. The automation functions and what they do are our biggest area of concern. The survey by McDonald et al. (2018) found that, for some automation functions, 20% of all drivers were unaware of whether or not the functions were available in their own cars. Then comes the basic question of which aspects of the driving task are supported by each automation function. Researchers at the University of Utah have compiled a list of automation functions and the variety of names assigned to them used by major car manufacturers today. They found that similar-sounding names were assigned to features that offered different functionalities creating an array of confusion, even among drivers who reported that they understood the automation well (Funkhouser, Tanner, & Drews, 2017). For example, a lane-keeping system might invite three different interpretations. In one interpretation, the car is able to automatically steer the car within a lane. In another interpretation, the system requires the driver to steer but then forcefully intervenes when the lane is inadvertently departed. In yet another interpretation, a warning is issued when the lane is departed, leaving the driver to evaluate and correct the situation. A survey reported in Stewart (2018) found that 70% of all drivers believe that autonomous cars are available for purchase today, whereas 11% of all drivers would consider taking a nap, reading a newspaper, or watching a movie when the highway pilot feature was engaged. This survey documents wild inaccuracies in drivers' pre-purchase understanding of how these cars work—placing a premium on training for those who will buy or operate them. We believe that drivers can and will operate these cars under mistaken assumptions about what functions they offer and that these confusions will be aggravated when drivers operate rental cars or cars borrowed from family and friends.

Engineering logic. We found many examples of mode complexities in the car manuals we read. At any given time, most of the automation functions could be in use (engaged), available for use (armed), or not available for use (off)—with a variety of conditions to determine which mode or modes the system could be in. Lane assist features can only be engaged when the vehicle has attained a minimum speed. Adaptive cruise control can only be engaged at a minimum speed but then disengages when the vehicle is decelerated through a different minimum speed. As memory for these details is likely to be poor, it is critical for the driver to monitor the annunciations appearing on the mode annunciator display. Like pilots, drivers may neglect to use these displays because they do not fully understand the implications of what is being displayed on them or when they are confused by the design of the display itself (McDonald et al., 2018; Resnick, 2016; Sarter, Mumaw, & Wickens, 2007). More concerning still are the *uncommanded* transitions between modes. In these situations, an automation mode can disengage itself or pass control to another mode that performs some but not all of its functions, while providing the operator with subtle feedback. We believe that these systems have grown too complex to be safely used without a basic understanding of these logic concepts. Most importantly, drivers need to understand that the mode displays, however well or poorly designed they may be, are the driver's view onto the workings of the automation: often the only transparency that the system provides. Endsley (2017) has provided an analysis of the functions available in one particular make and model vehicle. This sort of analysis should be available for every vehicle in the fleet.

Limitations. When we carefully examined the owner's manuals of several automated cars, we found many examples of limitations of the car automation that could significantly affect driver outcomes. Pedestrians under a specified height may not be detected by the pre-collision warning system. Functions available in one geographic region may not be available in others. Again, studies have already documented poor awareness of these system limitations among owners of automation-equipped cars (McDonald et al., 2018; Resnick, 2016).

Beyond static limitations such as minimum speeds and pedestrian heights, there are a great number of operational limitations—some of them unknown to anyone. When the adaptive cruise control is tracking a car in front and that car moves to the left or right, time is required to reassess the situation to see whether another car is ahead. If there is a stopped car ahead that the previous car just swerved to avoid, it is unlikely that the automation will be able to handle the situation. Similar limitations occur when rounding curves in which the adaptive cruise control can suddenly begin tracking a car ahead in a different lane rather than the car ahead in the same lane. Given that these automation systems operate using trained networks rather than simplistic rules, the behavior of the car in any given situation remains an empirical question: there is no way to know what the car will do other than to try it.

Driver struggles with these issues are already being documented. McGhee (2016) reports that 40% of respondents reported their vehicle had acted in a way that startled them or in a manner they did not expect. The survey by Loeb et al. (2018) shows that experience with automation surprises only increases drivers' desire for training.

Equipment. During our jet training, we were required to draw diagrams of aircraft systems from memory: complete with each significant component. We often criticized this requirement because we failed to see the point. Would we remember all of these details? What good would it do us if we did? The industry eventually arrived to the concept of *actionable* system knowledge—knowledge that could help us understand a system behavior that we hadn't seen before and guide us in choosing our next steps. Although we agree that diagramming the innards of an internal combustion engine holds little value for drivers, we find several examples of hardware installed in partially automated cars today that require actionable knowledge for their safe use. For example, adaptive cruise control relies on a single grill-mounted radar sensor and that even a drive through a mud puddle could incapacitate the system instantly. The camera used by the lane-keeping system can be similarly blocked by anything mounted on the roof such as a surfboard, skis, or perhaps most common: a mattress. The radar sensors mounted on the rear bumper are easily obscured by bumper stickers which could easily be affixed by a child. The current design seems to make the assumption that, in the case of a system that is inoperative, or that becomes inoperative while driving, the indications and alerts presented by the system will be sufficient to advise any driver that the system is no longer working. Our studies of how pilots respond to alerts and alarms leave us with little confidence that the design alone will serve to foolproof any system used by human drivers. We believe a basic understanding of how the workings of cameras, sensors, and other equipment can affect the way that the car responds to driver inputs is necessary. But while the idea of a "pre-trip walkaround" is a familiar one for pilots and commercial drivers, it is likely a new concept for drivers.

Knowledge About the Driver

Studies of driver behavior in the presence of car automation have already established a need to inform drivers about the likely effect of automation on their own behavior.

Maintaining awareness. Few pilots anticipated that a modern cockpit could surround them with every conceivable detail about their situation only to find that their awareness of it has been diminished. Driving studies have already replicated our findings in aviation: driver monitoring naturally decreases as the level of automation used increases (Carsten et al., 2012) and these attentional excursions have deleterious effects (Merat, Jamson, Lai, & Carsten, 2012). The findings that pilots tend to rely on warning systems as a primary means of monitoring their situation has also been demonstrated in cars (Ruscio, Ciceri, & Biassoni, 2015).

Studies of driver use of rearview cameras to detect pedestrians have revealed even more insidious problems with this type of safety feature. A large category of pedestrian crashes occurs when a child comes running "out of nowhere" toward the back of the vehicle (Fenton, Scaife, Meyers, Hansen, & Firth, 2005; Nhan, Rothman, Slater, & Howard, 2009). Of great concern is whether drivers would exclusively use the rearview camera and only scan the area directly behind the vehicle and skip the traditional head-checks that allow them to additionally scan the periphery. Indeed, one study found that rearview cameras only helped drivers avoid hitting stationary targets behind the car (Kidd, Hagoski, Tucker, & Chiang, 2015).

Admittedly, reminding drivers about the limits of their own attention-paying abilities does not enjoy a history of success (Casner, 2017). Nevertheless, gains made in military (Endsley & Robertson, 2000) and aviation applications (Flight Safety Foundation, 2014), together with the success of providing basic advice to drivers about automation (Barg-Walkow & Rogers, 2016), suggest that continued emphasis on the limitations of human attention is worthwhile.

Responding to abnormal events. Existing studies document drivers' steadfast belief that they will be able to look up, take the controls, and respond to situations that pop up (Xiong, Boyle, Moeckli, Dow, & Brown, 2012), a process known to human factors professionals to be more complicated and less reliable than most imagine. These mistaken beliefs negatively affect drivers when they use adaptive cruise control systems that allow them to choose a following distance behind the car being tracked in front of them. Studies show us that people are poor judges of their own reaction times during expected and unexpected events (Fitch, Blanco, Morgan, & Wharton, 2010; Taoka, 1989) as well as the time it takes them to choose a reaction when one must be chosen among several alternatives (Green & Von Gierke, 1984). The distinction between simple and choice reaction times is something that must be taught, even to students who seek out an understanding of these psychological nuances. Looking back to our discussion of driver awareness, the existing data show that driver takeover ability is further compromised in the absence of full attention (Merat et al., 2012).

Knowledge of the New Driving Task

Drivers need to know that they are responsible for the operation of the car, that the automation functions are tools at their disposal, and it is up to them to decide why, how, and when to use these tools. This sort of proactive thinking is different from using automation functions solely because they are installed in the car. We argue that these concepts too will need to be taught to drivers. Again, our experience with cockpit automation provides us with some insight about what form that knowledge might take.

Degani and Wiener (1994) proposed the idea that every pilot (and driver) should have an *automation philosophy*: personal reasons for using the automation features. Why would a driver want to use a lane-keeping assist and adaptive cruise control feature? Why have a computer perform these tasks instead of the driver? Consider the case of the driver who understands that "time-saving opportunities," such as highway lane changes, seldom yield anything other than additional risk for everyone

(Ellison & Greaves, 2015). Yet that driver still can't help responding to an opportunity to pass another car or to an aggressive move made by another driver. Using the highway pilot feature forces these reactions to pass through a planning and execution cycle—a lengthy sequence of disengaging, re-setting, and then re-engaging the automation. The work required to act on these irrational urges may render them no longer worth the trouble.

Degani and Wiener realized that every situation is different and that a single philosophy isn't enough to guide us in every situation that comes our way. So they proposed the idea of having *automation policies* that inform the automation philosophy in specific situations. For example, some drivers might feel that the highway pilot feature is a good tool to use when they are feeling drowsy. Why have a dozing driver at the controls when a robot can operate them with greater reliability? But the available studies tell us that using the automation naturally reduces driver attention. A drowsy driver needs engagement not disengagement. Driving when fatigued is a highly inopportune time to use this feature. So a good policy for using the highway pilot feature is to *avoid* using it when the driver's attention is already waning.

Realizing that what drivers do once the automation has been engaged is also important, Degani and Wiener proposed the concept of *automation procedures* to add details to the driver's use of the automation. While the most obvious automation procedure is to continue to pay attention to the road, a study by Loeb et al. (2018) uncovered more subtle details in what drivers are likely to do once the automation has been engaged. Loeb and her colleagues found that drivers who used the highway pilot feature tended to move their feet away from the brake pedal (and sometimes under the brake pedal). This is inadvisable because the autopilot could summon the driver back into the loop on a moment's notice and that having hands or feet away from the controls would only delay the driver response. Having an established place for hands and feet to rest while the autopilot is in use seems like a good procedure.

Automobile automation presents drivers with many options and most, if not all of them would benefit from advance thinking done outside the vehicle.

DISCUSSION

Ten years ago, flying a commercial jet airliner and driving a car were as unrelated as most any two activities could be. Today, leafing through a cockpit automation textbook (Casner, 2013), we find few concepts that we *wouldn't* want to pass on to drivers.

We have proposed three categories of knowledge that we believe every driver will need to possess to safely operate a car equipped with the automation systems that are being made available today. Table 1 summarizes our recommendations. We are aware that our first pass at a set of training standards is not complete nor is it free of elements that should ultimately be dropped. Other stakeholders should apply their expertise to what we have proposed and to refine our initial ideas. We invite individual manufacturers who deliver training to use our ideas as jumping off points to create their own training regimens.

We hear much talk about partially automated cars being a temporary "hold-me-over" as cars quickly step toward full autonomy along a continuum that has come to be known as the levels of automation (SAE International, 2018). There are a number of problems with this line of thinking. First, the timeline of this hold-me-over period is a topic of much debate. Although marketeers continue to promise next year or the year after (granting themselves frequent extensions), scientists counter with estimates of 75 years (Shladover, 2016). Regardless of whether we have autonomous cars in 2 years or a hundred, the partially automated cars that are sold this year won't receive upgrades that graduate them to higher levels of autonomy because autonomous cars use entirely different solutions. It won't be until these new cars are taken out of service that we will be free of any unaddressed human factors problems they introduce. Given that the average age of cars on the road is roughly 11 years, the partially autonomous cars that are released in coming years are going to be driven, en masse, for some time to come, regardless of whatever else happens.

A next important question that arises is: Who will deliver training to drivers? In a survey

TABLE 1: Minimum Driver Training Elements for Partially Automated Cars

Automation

- The automation functions offered by the car
- Which aspects of the driving task are assumed or assisted by each automation function
- How each automation function can be activated and deactivated by the driver
- How the car can automatically activate or deactivate each function (when applicable)
- How the driver can tell which functions are active at any given time
- The cameras and sensors upon which the automation functions depend
- How cameras and sensors can be rendered inoperative
- The importance of visually inspecting the hardware before driving
- How to recognize an inoperative camera or sensor
- Each automation function has situations in which it will not work or stop working: some situations are listed in the driver manual and some are unknown even to the engineers who designed the system
- The behavior of the car can change with a software update

Drivers

- Details about the car are easily forgotten or misremembered: even if initially well-learned
- Humans tend to overestimate their ability to pay attention and notice out-of-the-ordinary events
- Experienced pilots and drivers routinely miss things that happen right before their eyes
- Vigilance drops off quickly (~15–20 min) because of fatigue and mind wandering
- Humans are bad at noticing changes in a scene after having looked away for a few seconds
- Vigilance tends to wane when we delegate control to an automated system
- Drowsiness becomes more of a factor when control is delegated to automation
- When an alerting system is available to detect and call out threats, even the best-trained humans tend to rely on the alerting system and pay less attention to the out-the-window scene
- Pilots and drivers alike tend to overestimate their ability to take over control and respond to out-of-the-ordinary events
- When trouble pops up, attention becomes narrowed, alternatives are overlooked, reactions are delayed, time is lost. These are all essentially human traits. We are not robots

The Driving Task

- The car cannot drive itself, for any length of time
- The driver will likely be held responsible for whatever happens
- The automated features are tools available at your service, should you decide to use them
- Automation features should not be used because they exist or because of their cost
- It is useful to decide when are good and bad times to use the automation—before getting in the car

conducted by the State Farm insurance company, 52% of drivers stated a preference for being taught about their new car at the dealership at which they purchased the vehicle (Mullen, 2017). Fortunately, dealership training is already being made available. But Abraham, McAnulty, Mehler, and Reimer (2017) examined the training provided by dealerships at the point of sale and concluded that the quality of the training varied greatly. Having a standard set of knowledge and proficiency elements to work from, one that enjoyed inputs and refinements from all stakeholders could help make these training episodes more consistent.

High school driver education has gradually disappeared over the past decades, the decision to drop the program often justified by studies that show how current methods of driver training have little effect on the subsequent safety records of those who receive the training (Peck, 2011). That is, teens who completed driver training didn't seem to crash any less often than teens who didn't take the class. But looking at the historical crash statistics, most teen-involved crashes were attributable to aggressive driving and/or the use of alcohol. Peck (2011) argues that driver training that aims to simply provide more hours on the road may not be the right way to transfer knowledge or address driver attitudes toward these risky behaviors. But the risks presented by car automation may be different in that guided instruction, possibly even conducted in simulators, may give students a chance to gain experience that would be better received. It doesn't take long to witness an automated system do something unexpected and give one the impression that their attention is still needed. We argue that car automation and perhaps the surge in smartphone usage while driving are sufficient reasons to consider bringing back this part of secondary education.

In the aviation industry, we often say that good pilots are not trained, they are molded over time. Introductory training only plays a small part in the process of becoming a safe and proficient pilot. The learning that follows initial training happens in the aircraft and in the presence of another more experienced pilot. Granted, automobile trips generally only have one driver but there seem to be opportunities to use in-car automation to further inform and shape the behavior of drivers and to help drivers retain what they have previously learned (Casner, Heraldez, & Jones, 2006). We might study the interaction between flight crew members who benefit each other yet somehow avoid the "backseat driver" phenomenon that would be unwelcome in any vehicle. What timely comments and gestures are made by good co-pilots? Could those be mimicked by an in-car automation system? In addition to alerting drivers that they are about to hit something, perhaps the system could identify developing threats (e.g., a cyclist ahead who is not wearing reflective gear at night and might be hard to spot). The use of localized traffic and pedestrian data is another idea. Crash data are becoming more widely available. Troublesome roads and intersections could be pointed out to drivers with the suggestion of turning off the automation. Statistics on mid-block pedestrian traffic (jaywalkers) could be relayed to drivers in real time based on location.

Another important question is to what extent training should be focused on the design details of a particular make and model car and to what extent it should be aimed at teaching drivers the underlying concepts of how most any car automation system works. After 30 years of experience with aviation automation, we have strong opinions on this matter. Of course, manufacturers will use the knobs, dials, and details of their own designs as manipulatives when introducing drivers to their own cars: it would make little sense to do it any other way. But we remind manufacturers that there is a higher purpose here. Drivers will borrow and rent cars made by any number of manufacturers. What will help protect the short-term driver who was just tossed the keys to a car made by Manufacturer X? Good training provided by Manufacturer Y—the company who designed that driver's own car. Training that goes beyond knobology and button pushing and other cosmetic details. Training that conveys underlying concepts that are known to help drivers to more smoothly transfer what they know to cars that look different on the surface but are fundamentally similar in the way they work. The effectiveness of taking the time to provide richer, more conceptual knowledge has already been demonstrated in automation-equipped airplanes and the results are striking (Casner, 2005). We note that, when cars are rented and borrowed, a sort of prisoner's dilemma is created. Manufacturer X remains with fingers crossed that Manufacturer Y took the time to ensure that their customers would be safe when borrowing Manufacturer X's car—and vice versa.

Decades ago, the airline industry discombobulated pilots and human factors professionals alike by introducing revolutionary automation in airplanes flown with passengers in the back—with humble amounts of prior testing. Today, airline crashes are at a historic low—following a concerted effort by human factors professionals to raise awareness of how humans and automation systems can and must work together as a team, with an understanding of the strengths and limitations of each reflected in the other. Can history repeat itself in semi-automated cars?

ORCID ID

Stephen M. Casner https://orcid.org/0000-0003-2435-0331

REFERENCES

Abraham, H., McAnulty, H., Mehler, B., & Reimer, B. (2017). Case study of today's automotive dealerships: Introduction and delivery of advanced driver assistance systems. *Transportation Research Record*, *2660*, 7–14.

Barg-Walkow, L. H., & Rogers, W. A. (2016). The effect of incorrect reliability information on expectations, perceptions, and use of automation. *Human Factors*, *58*, 242–260.

Barry, K. (2018, May 30). Drivers more likely to use their phones when cruise control is on, study finds. *Consumer Reports*. Retrieved from https://www.consumerreports.org/car-safety/drivers-more-likely-to-use-phones-when-cruise-control-is-on-mit-study/

Boeing. (2017). *Statistical summary of commercial jet airplane accidents: Worldwide operations, 1959-2016*. Seattle, WA. Retrieved from https://bit.ly/2NwYiUa

Breen, B. C. (1999, January). Controlled flight into terrain and the enhanced ground proximity warning system. *IEEE Aerospace and Electronic Systems Magazine*, *14*, 19–24.

Breznitz, S. (1984). *Cry wolf: The psychology of false alarms*. Mahwah, NJ: Lawrence Erlbaum.

Bulfer, B. (2018). *FMC user's guide*. Merced, CA: Leading Edge.

Carroll, J. M., & Rossen, M. B. (1987). Paradox of the active user. In J. M. Carroll (Ed.), *Interfacing thought: Cognitive aspects of human-computer interaction* (pp. 80–111). Cambridge, MA, MIT Press.

Carsten, O., Lai, F., Barnard, Y., Jamson, A. H., & Merat, N. (2012). Control task substitution in semiautomated driving: Does it matter which aspects are automated? *Human Factors*, *54*, 747–761.

Casner, S. (2017, December 18). Why we can't stop texting and driving. *Time*. Retrieved from http://time.com/5059457/stop-texting-and-driving/

Casner, S. M. (2003). *Teaching cockpit automation in the classroom* (NASA Technical Memorandum 2003-211865). Moffett Field, CA: NASA Ames Research Center.

Casner, S. M. (2005). Transfer of learning between a small technically advanced aircraft and a commercial jet transport simulator. *International Journal of Applied Aviation Studies*, *5*, 307–319.

Casner, S. M. (2013). *The pilot's guide to the airline cockpit* (2nd ed.). Newcastle, WA: Aviation Supplies & Academics.

Casner, S. M., Geven, R. W., & Williams, K. T. (2013). The effectiveness of pilot training for abnormal events. *Human Factors*, *55*, 477–485.

Casner, S. M., Heraldez, D., & Jones, K. M. (2006). Retention of aeronautical knowledge. *International Journal of Applied Aviation Studies*, *6*, 71–97.

Consumer Reports. (2018, July 24). *Cars with advanced safety systems: Detailed list of cars with features that help drivers avoid or mitigate collisions*. Retrieved from https://goo.gl/36v7wt

Degani, A., & Wiener, E. L. (1994). Philosophy, policies, procedures, and practices: The four Ps of flight deck operations. In N. Johnston, N. McDonald, & R. Fuller (Eds.), *Aviation Psychology in Practice* (pp. 44–76). New York: Routledge.

Ellison, A. B., & Greaves, S. P. (2015). Speeding in urban environments: Are the time savings worth the risk? *Accident Analysis & Prevention*, *85*, 239–247.

Endsley, M. R. (2017). Autonomous driving systems: A preliminary naturalistic study of the Tesla Model S. *Journal of Cognitive Engineering and Decision Making*, *11*, 225–238.

Endsley, M. R., & Kiris, E. O. (1995). The out-of-the-loop performance problem and level of control in automation. *Human Factors*, *37*, 381–394.

Endsley, M. R., & Robertson, M. M. (2000). Training for situation awareness. In M. R. Endsley & D. J. Garland (Eds.), *Situation awareness analysis and measurement* (pp. 349–366). Mahwah, NJ: Lawrence Erlbaum.

Feary, M., McCrobie, D., Alkin, M., Sherry, L., Polson, P., & Palmer, E. (1998). *Aiding vertical guidance understanding* (NASA/TM-1998-112217). Moffett Field, CA: Ames Research Center.

Federal Aviation Administration. (2011). *Introduction to TCAS II: Version 7.1*. Washington, DC: Author. Retrieved from https://bit.ly/2J7Ke1C

Fenton, S. J., Scaife, E. R., Meyers, R. L., Hansen, K. W., & Firth, S. D. (2005). The prevalence of driveway back-over injuries in the era of sports utility vehicles. *Journal of Pediatric Surgery*, *40*, 1964–1968.

Fitch, G. M., Blanco, M., Morgan, J. F., & Wharton, A. E. (2010). Driver braking performance to surprise and expected events. In *Proceedings of the Human Factors and Ergonomics Society Annual Meeting*, *54*(24), 2075–2080. https://doi.org/10.1177/154193121005402412

Flight Safety Foundation. (2014). *A practical guide for improving flight path monitoring. Final report of the active pilot monitoring working group*. Retrieved from: https://flightsafety.org/files/flightpath/EPMG.pdf

Funkhouser, K., Tanner, E., & Drews, F. (2017, June). *Know it by name: Human factors of ADAS design*. Poster presented at the 2017 Autonomous Vehicles Symposium, San Francisco, CA.

Green, D. M., & Von Gierke, S. M. (1984). Visual and auditory choice reaction times. *Acta Psychologica*, *55*, 231–247.

Hergeth, S., Lorenz, L., Vilimek, R., & Krems, J. F. (2016). Keep your scanners peeled: Gaze behavior as a measure of automation trust during highly automated driving. *Human Factors*, *58*, 509–519.

Hutchins, E. (2007). Measuring change in pilots' conceptual understanding of autoflight. In *Proceedings of the 14th International Symposium on Aviation Psychology*, (pp. 281-286). Dayton, Ohio: Wright State University.

Kanki, B. G., Anca, J., & Helmreich, R. L. (2010). *Crew resource management*. Cambridge, MA: Academic Press.

Kidd, D. G., Hagoski, B. K., Tucker, T. G., & Chiang, D. P. (2015). The effectiveness of a rearview camera and parking sensor system alone and combined for preventing a collision with an unexpected stationary or moving object. *Human Factors*, *57*, 689–700.

Körber, M., Baseler, E., & Bengler, K. (2018). Introduction matters: Manipulating trust in automation and reliance in automated driving. *Applied Ergonomics*, *66*, 18–31.

Ledesma, R. D., Montes, S. A., Poó, F. M., & López-Ramón, F. (2015). Measuring individual differences in driver inattention: Further validation of the attention-related driving errors scale. *Human Factors*, *57*, 193–207.

Loeb, H., Belwadi, A., Shaikh, S., Shen, M., & Ward McIntosh, C. (2018, July 11). *Emergency autonomous to manual takeover in a driving simulator: Teens vs. adults, males vs. females*. Poster presented at 2018 Autonomous Vehicles Symposium, San Francisco, CA.

McDonald, A., Carney, C., & McGhee, D. V. (2018, September). *Vehicle owners' experiences with and reactions to advanced driver assistance systems*. Washington, DC: American Automobile Association Foundation for Traffic Safety.

McGhee, D. (2016, October 27). *Driver perceptions of ADAS (My does what survey). Presentation given at National Transportation Safety Board and National Safety Council meeting—Reaching Zero Crashes: A Dialogue on the Role of Advanced*

Driver Assistance Systems (ADAS). Retrieved from https://bit.ly/2uM0wI6

Merat, N., Jamson, A. H., Lai, F. C., & Carsten, O. (2012). Highly automated driving, secondary task performance, and driver state. *Human Factors*, *54*, 762–771.

Miller, E. E., & Boyle, L. N. (2017). Driver adaptation to lane keeping assistance systems: Do drivers become less vigilant? In *Proceedings of the 61st Annual Meeting of the Human Factors and Ergonomics Society* (pp. 1934–1938). Washington, DC: Human Factors & Ergonomics Society.

Ministry of Transport. (1996). *Aircraft accident investigation report* (China Airlines, Airbus Industry A300B4-622R, B1816, Nagoya Airport, April 26, 1994). Nagoya, Japan: Aircraft Accident Investigation Commission, Ministry of Transport.

Mullen, C. (2017). *Presentation given at Reaching zero crashes: A dialogue on the role of current advanced driver assistance systems*. Washington, DC: National Transportation Safety Board and the National Safety Council.

National Transportation Safety Board. (1989). *Aircraft accident report: Delta Air Lines, Inc. Boeing 727-232, N473DA, Dallas-Forth Worth International Airport, Texas*. *August 31, 1988* (Report NTSB/AAR-89/04). Washington, DC: Author.

National Transportation Safety Board. (2014). *Descent below visual glide path and impact with seawall Asiana Airlines Flight 214, Boeing 777-200ER, HL7742, San Francisco, CA*, *July 6, 2013* (Report NTSB/AAR-14/01). Washington, DC: Author.

Naujoks, F., Kiesel, A., & Neukum, A. (2016). Cooperative warning systems: The impact of false and unnecessary alarms on driver compliance. *Accident Prevention & Analysis*, *97*, 162–175.

Nhan, C., Rothman, L., Slater, M., & Howard, A. (2009). Back-over collisions in child pedestrians from the Canadian Hospitals Injury Reporting and Prevention Program. *Traffic Injury Prevention*, *10*, 350–353.

Nolan, M. S. (2010). *Fundamentals of air traffic control*, *5th edition*. Clifton Park, NY: Delmar Cengage Learning.

Parasuraman, R., & Riley, V. (1997). Humans and automation: Use, misuse, disuse, abuse. *Human Factors*, *39*, 230–253.

Peck, R. C. (2011). Do driver training programs reduce crashes and traffic violations?—A critical examination of the literature. *International Association of Traffic and Safety Sciences Research*, *34*, 63–71.

Resnick, M. (2016, June 29). *What is the automation doing? Mode awareness problems catch Tesla by surprise. Ergonomics in design*. Retrieved from https://bit.ly/2zPn7p9

Rettig, M. (1991). Nobody reads documentation. *Communications of the ACM*, *34*(7), 19–24.

Ruscio, D., Ciceri, M. R., & Biassoni, F. (2015). How does a collision warning system shape driver's brake response time? The influence of expectancy and automation complacency on real-life emergency braking. *Accident Analysis & Prevention*, *77*, 72–81.

SAE International. (2018). *Taxonomy and definitions for terms related to driving automation systems for on-road motor vehicles* (Report J3016_201806). Warrendale, PA: Society of Automotive Engineers International.

Sarter, N. D., Mumaw, R. J., & Wickens, C. D. (2007). Pilots' monitoring strategies and performance on automated flight decks: An empirical study combining behavioral and eye-tracking data. *Human Factors*, *49*, 347–357.

Sarter, N. D., & Woods, D. D. (1992). Pilot interaction with cockpit automation: Operational experiences with the flight management system. *International Journal of Aviation Psychology*, *2*, 303–321.

Sarter, N. D., & Woods, D. D. (1995). How in the world did I ever get into that mode? Mode error and awareness in supervisory control. *Human Factors*, *37*, 5–19.

Sarter, N. D., & Woods, D. D. (1997). Team play with a powerful and independent agent: Operational experiences and automation surprises on the Airbus A-320. *Human Factors*, *39*, 553–569.

Sarter, N. D., Woods, D. D., & Billings, C. E. (1997). Automation surprises. In G. Salvendy (Ed.), *Handbook of human factors & ergonomics* (2nd ed., pp. 1926–1943). New York, NY: John Wiley.

Shladover, S. E. (2016). The truth about "self-driving" cars. *Scientific American*, *314*, 52–57.

Stewart, J. (2018, November 18). Drivers wildly overestimate what "semiautonomous" cars can do. *Wired Magazine*. Retrieved from https://bit.ly/2Ot6xFs

Taoka, G. T. (1989). Brake reaction times of unalerted drivers. *ITE Journal*, *59*, 19–21.

Thackray, R. I. (1965). Correlates of reaction time to startle. *Human Factors*, *7*, 74–80.

Warrick, M. J., Kibler, A. W., & Topmiller, D. A. (1965). Response time to unexpected stimuli. *Human Factors*, *7*, 81–86.

Watzenig, D., & Horn, M. (2017). *Automated driving: Safer and more efficient future driving*. New York, NY: Springer.

Wiener, E. L., & Curry, R. (1980). *Flight-deck automation: Promises and problems* (NASA Technical Memorandum 81206). Moffett Field, CA: NASA Ames Research Center.

Xiong, H., Boyle, L. N., Moeckli, J., Dow, B. R., & Brown, T. L. (2012). Use patterns among early adopters of adaptive cruise control. *Human Factors*, *54*, 722–733.

Stephen M. Casner (stephen.casner@nasa.gov) is a research psychologist at NASA Ames Research Center, Moffett Field, CA. He holds an FAA Airline Transport Pilot certificate, type ratings in the Boeing 737 and Airbus A-320 aircraft, and a Gold Seal Certified Flight Instructor certificate.

Edwin L. Hutchins (ehutchins@ucsd.edu) is an emeritus professor of cognitive science at the University of California, San Diego. He holds an FAA Commercial Pilot certificate and type ratings in the Douglas DC-3 and the Cessna Citation 500-series jet aircraft.