INFORMATION SYSTEMS EDUCATION JOURNAL

Volume 21, No. 3
July 2023
ISSN: 1545-679X

In this issue:

This issue focuses on three areas of cybersecurity (two in teaching case study form), a chatbot for grad advising, and two papers on student competencies as they relate to curriculum. The chatbot leverages AI to provide automated chat responses for grad student advising questions, while the competency papers address integration of the IS2020 Elective competencies into the curriculum, along with a study of digital competencies of the incoming student class in a first-year experience course. Our first cybersecurity paper includes a study of cybersecurity defenses in health care, which might readily be used in cybersecurity or health care informatics courses. The other two papers are case studies – one focusing on risk management in manufacturing, and the other on the risk elements of autonomous vehicles.

iscap Information Systems & Computing Academic Professionals

The **Information Systems Education Journal** (ISEDJ) is a double-blind peer-reviewed academic journal published by **ISCAP** (Information Systems and Computing Academic Professionals). Publishing frequency is five times per year. The first year of publication was 2003.

ISEDJ is published online (https://isedj.org). Our sister publication, the Proceedings of EDSIGCON (https://proc.iscap.info) features all papers, abstracts, panels, workshops, and presentations from the conference.

The journal acceptance review process involves a minimum of three double-blind peer reviews, where both the reviewer is not aware of the identities of the authors and the authors are not aware of the identities of the reviewers. The initial reviews happen before the ISCAP conference. All papers, whether award-winners or not, are invited to resubmit for journal consideration after applying feedback from the Conference presentation. Award winning papers are assured of a publication slot; however, all re-submitted papers including award winners are subjected to a second round of three blind peer reviews to improve quality and make final accept/reject decisions. Those papers that are deemed of sufficient quality are accepted for publication in the ISEDJ journal. Currently the target acceptance rate for the journal is under 36%.

Information Systems Education Journal is pleased to be listed in the Cabell's Directory of Publishing Opportunities in Educational Technology and Library Science, in both the electronic and printed editions. Questions should be addressed to the editor at editor@isedj.org or the publisher at publisher@isedj.org. Special thanks to members of ISCAP who perform the editorial and review processes for ISEDJ.

INFORMATION SYSTEMS EDUCATION JOURNAL

Teaching Case

Yours, mine and ours: Risk assignments, management, and tradeoffs on the road to driverless vehicles

Paul D. Witman
pwitman@callutheran.edu

Jim Prior
jr2dg.consulting@gmail.com

School of Management
California Lutheran University
Thousand Oaks, CA, USA

Scott Mackelprang
scottmackelprangemail@gmail.com

Abstract

This case study looks both backwards and forwards at real and potential incidents and risks created by the use of various levels of self-driving vehicles. The case provides background on autonomous vehicle technology, the legal and risk management frameworks involved, and a variety of scenarios for students to consider. The scenarios provide the foundation for discussion of autonomous vehicle introductions and operations, as well as considerations for how other new technologies may be launched.

Keywords: autonomous vehicles, technology rollout, risk management

1. INTRODUCTION

It is time to go to meet friends for dinner. You open the door of your car, get in, and tell the car where to take you. There are no controls other than the screen which shows the route you will be taking. The car drives up to the door of the restaurant, drops you off, and then finds a place to park. After dinner, you use your phone to ask the car to come pick you up and take you back to your home. When you are not actively using your car, you can lend it out to others to earn some extra cash.

Many people will reach an age or condition when they can no longer drive safely. This might be due to reaction time, vision, confidence, or physical limitations. It would be wonderful to still have the same freedom of movement that is provided by driving - which may be possible if we can get technology to do some of the work for us. But such vehicles are not here yet - and current versions still have quirks, like the Waymo vehicles that wrongly routed dozens of trips into a dead-end *cul-de-sac* in San Francisco, disrupting traffic for those residents (CBS News Bay Area, 2021).

The spectrum of automation that can be implemented in driver-assisted vehicles can best be understood using the Society of Automotive Engineers (SAE) classification system, illustrated in Figure 1. This model defines levels 0, 1, and 2 as Advanced Driver Assistance Systems (ADAS), which require a driver. It also defines levels 3, 4,

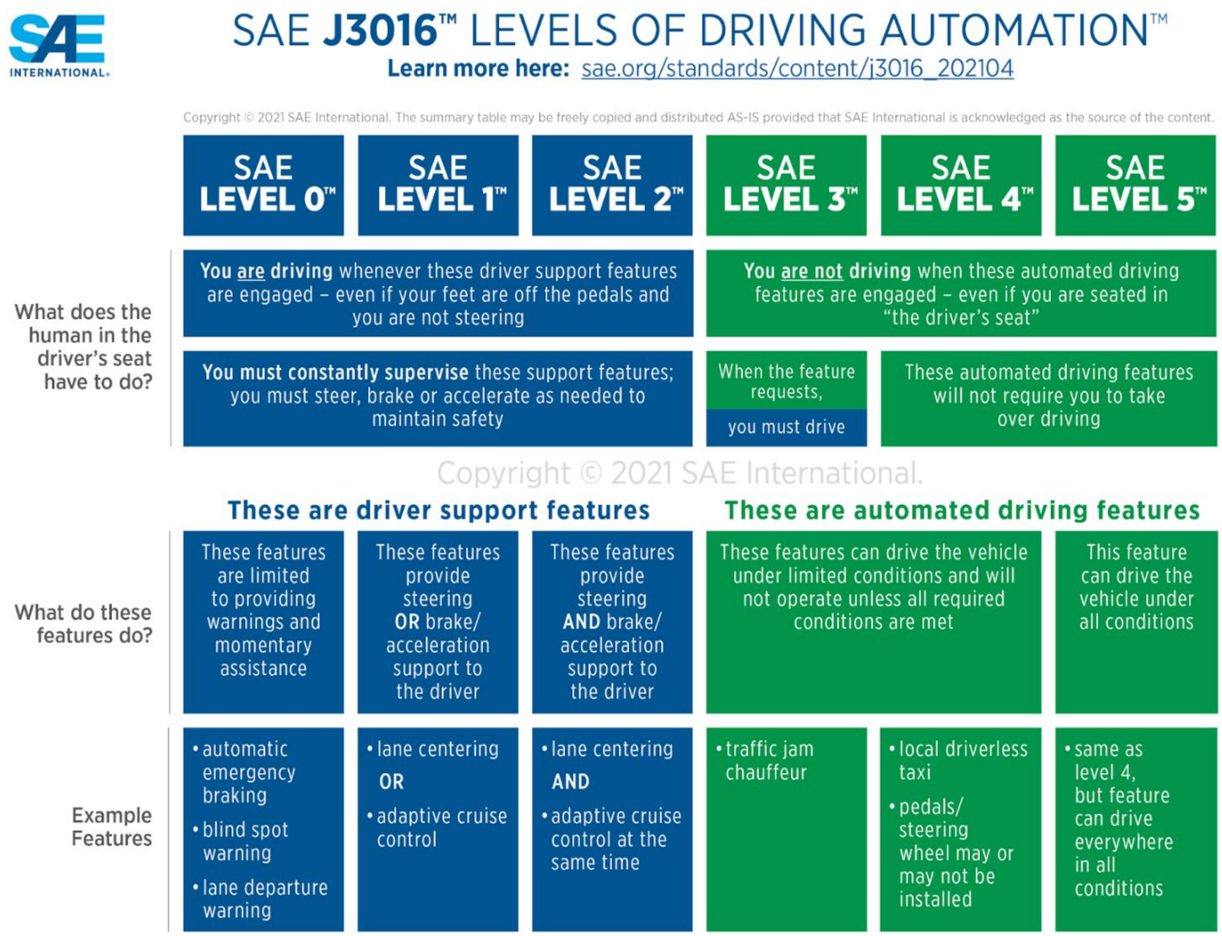

Figure 1: SAE Levels of Driving Automation (Society of Automotive Engineers International, 2021), used with permission

and 5 as Automated Driving Systems (ADS), which allow complete control of the vehicle by the technology under a variety of limited conditions, ultimately leading to a vehicle that requires no human driver (Society of Automotive Engineers International, 2021).

Many ADAS features are available in cars today that provide features such as adaptive cruise control, forward collision warning, automatic emergency braking, lane departure warning, lane-keeping assistance, and blind-spot warning as well as more advanced features such as lane-centering technology, lane-changing assistance, and self-steering systems. Drivers may already be developing dependencies on these capabilities and adjusting their driving habits because of them. However, there are currently no Automated Driving Systems available for consumer purchase in the U.S., though some are being used in tests for taxis and other such commercial operations (U.S. Department of Transportation, 2021).

There are many potential benefits from current and future levels of driving automation. Proponents expect to achieve increased safety for drivers and passengers along with other motorists and pedestrians, as well as increased mobility via expanded transportation options. They also expect economic and social benefits from a reduction in car crashes, less need for parking spaces/lots, increases in automated ride sharing reducing environmental impact, and efficiency and convenience with reduced traffic congestion (Litman, 2017).

Of course, the automation of driving brings challenges as well: legal liability and insurance systems to accommodate and cover accidents of autonomous and semi-autonomous vehicles, the need for new state and federal regulations around risk assignments, and technology limitations that make fog and rain, or poor road markings more difficult to manage safely. Risks stemming from malicious hacking of autonomous vehicles can be anticipated, as can the potential for privacy issues as the vehicles coordinate their movements among themselves and share their observations about road conditions. It is also reasonable to expect people to use vehicles in ways that are good for them, while less good for society – e.g., sending their car away from a downtown location to find cheaper parking, saving money but creating traffic. All of these must be appropriately weighed against the economic and social value of entrepreneurship and innovation, lest we over-regulate.

Autonomous vehicles will change how accident liability is decided in the courts. Over time, individual court cases across the spectrum of SAE levels of driving automation will be treated as precedents contributing to autonomous vehicle case law. Judges and juries will have to contend with potentially complex technical issues. Was the accident due to a failure of a collision avoidance subsystem manufactured by a third party? Or a software failure of some kind? Did the driver fail in their ultimate responsibility? Or did the vehicle's automation fail them? Was the system not maintained properly? What do the current and future laws say about liability in those situations?

Not surprisingly, it takes a lot of technology to try to replicate and improve on the human driver. Vehicles need to know where they are in terms of geographical location, they need to know what is around them, how to react to those things, and how to interact with them appropriately. Many types of sensors are involved, as highlighted in Figure 2, below. Sensors may include things like visible light cameras, short- and long-range radio-based distance tracking (radar), ultrasound, and light-based distance tracking (called LIDAR). Data from all of these sensors is then integrated in high-performance computing devices that determine how to change speed and direction as required (Wendt & Cook, 2018). In addition to the sensors, vehicles need Global Positioning System (GPS) or Global Navigation Satellite System (GNSS) data from satellites, detailed maps (more detailed than are publicly available), and map updates based on recent data from other vehicles which may have discovered things like lane closures, etc.

To make sure that these systems, whether ADAS or automated driving, work properly, they go through many levels of testing by the manufacturer. Eventually, the vehicles need to be tested in a "real-world" environment – the actual roads and traffic conditions that the vehicles need to operate in. This naturally introduces new risks to occupants of automated vehicles: an error

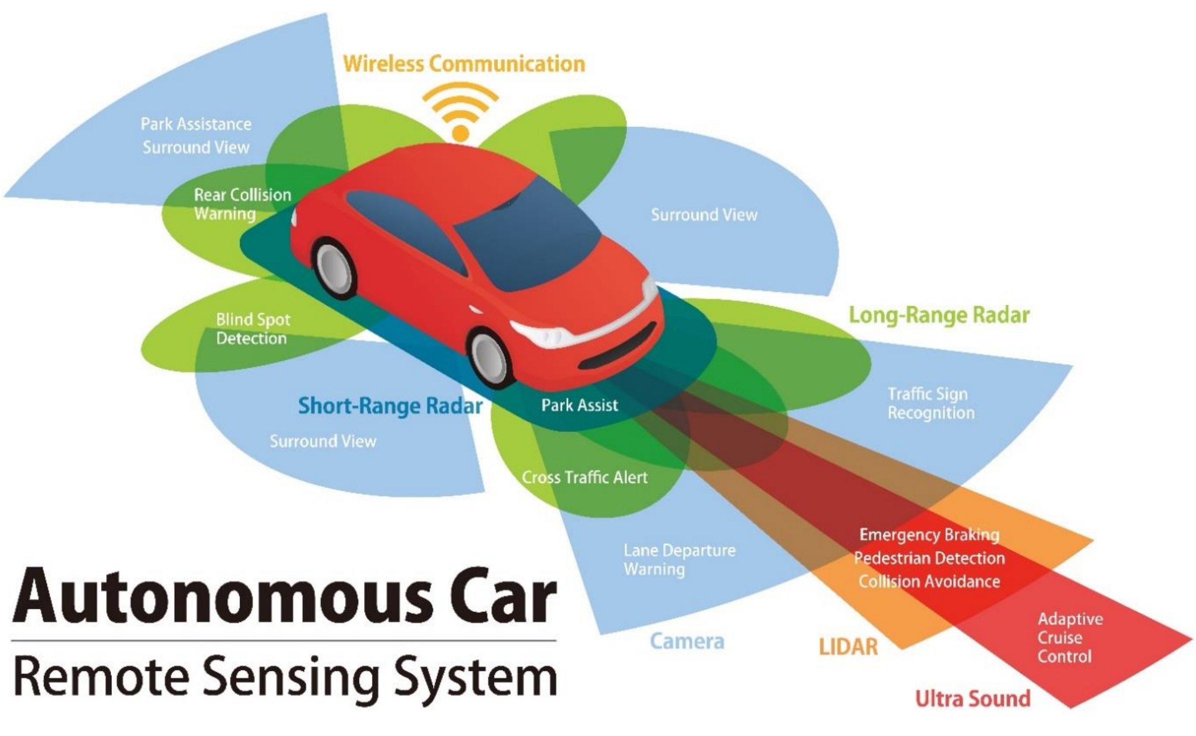

Figure 2: Sensor components of autonomous vehicles (metamorworks, 2022)

could injure passengers and drivers as well as damage the vehicle. It also introduces new risks for other nearby drivers and pedestrians who might be injured if automated vehicles make an error.

Risk is a concept that is not unique to business or technology - it is fundamental to many aspects of everyday life. Risk refers to the probability that a certain type of event will occur. Related to that concept of risk is the magnitude of impact that might occur should the event occur. Risks can result in positive or negative outcomes, though we most commonly focus on risk as being related to negative outcomes - injuries, damage, cost, and the like.

There are several fundamental approaches to managing risk. One can avoid the risk, transfer it to someone else, reduce it to acceptable levels or simply accept it.

> *Avoidance*: If a risk is determined to be too high, then you avoid the activity that creates the risk.
>
> *Transfer*: In many cases, the risk can be transferred to another party.
>
> *Reduction*: Risk reduction is a common step for processes or activities that cannot be avoided, and where the risk cannot be transferred to another party.
>
> *Acceptance*: In some cases, the best option is to accept a risk (Ahmed, 2017). If steps have been taken (as in the above strategies) to reduce or mitigate the risk, the remaining level of accepted risk is known as "residual risk."

The dynamics of the risk management being exercised in the advertising, sale, regulation, operation, and enforcement of autonomous vehicles are currently fluid and not well understood by the public. This will, undoubtedly, lead to challenges for the industry.

Ethical choices, for the manufacturer, the user, and other parties, are generally viewed through one of several different "lenses." These lenses include (among others) the fundamental "rights" of affected parties, a view of what represents *justice* in the broader context of merit and need, a utilitarian view that balances the resulting good vs. harm across all parties involved, and a "common good" approach that is similar to the utilitarian approach, but recognizes a mutual concern for all parties involved directly or indirectly (Markkula Center for Applied Ethics, 2021).
For the past several decades technology has advanced more rapidly than the societal and governmental structures to manage and coordinate them (Malan, 2018). For example, e-cigarettes and social media, which use advanced technology in dynamic economic and operational environments, both confronted society in advance of the regulatory structures to govern them. Manufacturers and users of autonomous vehicles are now beginning to face the challenges of managing their legal risks and liabilities without well-defined legal structures. Such risks can impact the rollout or adoption of autonomous vehicles. In at least one case, release of an otherwise-ready technology, Audi's Level 3 self-driving for traffic jams, was deferred largely due to lack of a legal framework to manage its risk and potential liability (Szymkowski, 2020).

2. DISCUSSION SCENARIOS

The following are several scenarios to consider as to how the introduction of this new technology has been or could be managed. The goal of these scenarios is to spark your thinking about the various factors that impact the deployment of new technologies. These factors might include issues of risk, liability, insurance, legal structures, ethical considerations, system testing and more. Each scenario includes a real or hypothetical situation, and the technology and other decisions and factors that affect it, and it asks you to consider the tradeoffs and decisions involved.

Risks, Broadly Considered

Mary heads off to work in her car, driving herself as usual. Two blocks away, Sam leaves about the same time heading to a destination close to yours, but lets his car choose the route, navigate, and control the vehicle's operation. Mary encounters Sam's car on the way, to her right at a four-way stop, arriving just ahead of Sam's car. Mary thinks she has the right of way, but can't make eye contact with Sam, even after honking. She waits, as does Sam's car – neither one immediately entering the intersection. When Mary decides after a few seconds to go ahead, she sees Sam's car start to do the same. After another false-start, Mary finally proceeds across the intersection, followed by Sam's car. Mary has just encountered an unexpected risk (not getting the usual cues from the other driver), which she has coped with by paying attention, carrying insurance, etc. All drivers take a number of measures to reduce the risk of adverse events happening, and to reduce the damage if they do, or to transfer some of the risk to others.

While there are many potential benefits to autonomous vehicles, benefits always come with some new risks – as with every new technology. Drivers may be more inclined to risk "driving" while intoxicated, since the car will "do most of the driving." Drivers have already, many times, bypassed "alert driver" detection systems (e.g., Dunn et al., 2021; Lambert, 2021), which defeats those controls.

These new risks also extend to other drivers, pedestrians, first-responders, cyclists, and road workers. Even though a self-driving car may be statistically safer than the average human driver, there have still been cases of those vehicles behaving erratically, in ways that are unexpected by their human-driven counterparts, or that require the ADAS driver to take control (Goodall, 2021; Smiley, 2022). Either of these may result in damage, injury, or death to the ADAS driver and to all those around them.

Among the many tradeoffs that arise with autonomous vehicles, there is a phenomenon called automation complacency (Azuma et al., 2022). When humans are responsible for operating a device that largely operates independently and successfully, it is common for the operator to lose focus, and to assume the device will always do the right thing, even when their own safety and the safety of those around them depend on their ability to retake control of the vehicle quickly.

Questions

- When a car manufacturer sells you a self-driving car, and you sign a contract saying that you understand that you are responsible for its safe operation, which risk management approach is the manufacturer using?
- Is that the "right" approach? How would you justify your answer?
- What are the ethical implications for the driver/owner, as well as for the public?
- If you as an autonomous vehicle owner want to reduce the risks associated with autonomous driving, what steps and decisions could you make to reduce the risk that your vehicle injures anyone, or damages property?

Cybersecurity Risks of Autonomous Vehicles

In 2015, two self-described hackers demonstrated an attack on a Chrysler Jeep Cherokee. In the attack, they used an insecure channel through the entertainment system of the vehicle to take control of air conditioning, radio volume, the picture on the dashboard screen, driving speed, and windshield wipers. The hackers were able to do this remotely, via a wireless Internet connection, and required no direct physical access to the Jeep's systems (Greenberg, 2016).

Now think ahead to a Level 4 or 5 autonomous vehicle, with no driver controls available. The vehicle navigates based on input from GPS satellites, maps downloaded from a central server, and the measurements of sensors in the vehicle, along with (perhaps) direct communication with other nearby vehicles.

Knowing that hackers are motivated by a wide range of goals (money, fun, power, bragging rights, etc.), hackers might attempt to conduct random attacks on cars, or might target a specific vehicle or vehicles to try to disrupt, frighten, or injure that vehicle's rider. Hackers might try to take remote control of the vehicle, as in the story above, or might try to disrupt its sensors in some fashion.

Questions

- Identify the navigating and driving control inputs that might be attacked, and some potential approaches to prevent or mitigate those attacks.
- How might a hacker attack a vehicle in a way that could generate income for the hacker?
- What do manufacturers need to do to mitigate the risk of cyberattacks?
- What privacy considerations arise when society is being driven around in autonomous vehicles?
- Should there be identification and authorization requirements for operating an autonomous vehicle? How might they be implemented to help prevent large scale vehicle theft in the event of a centralized vehicle controller breach? How might such control measures impact the ability of owners to lend their car to a friend?

Marketing vs. Rigorous Definitions of "Self-Driving"

Melissa has gone out to buy a new car and is considering buying a Tesla. She notes that it has a feature called "Autopilot." She's also considering a Honda, which has a feature called "Traffic Jam Pilot," and a Mercedes Benz, who call their feature "Drive Pilot." Melissa is at a bit of a loss to understand exactly what these features do, without looking deeply into the sales material.

What do you think Melissa sees when she reads the word *autopilot* on its own? Perhaps an image of an airline pilot sitting back, letting a system control an airplane in flight?

What about the Tesla feature called "Autopilot"? Same thing? Or is it a new concept since it is related to a car? What about "Drive Pilot" or "Traffic Jam Pilot"? How do those differ, if at all, from "Autopilot"?

These companies may be branding their conditional automated driving systems in this manner, hoping that you will think their systems are like those used in airplanes.

But are they? Are the circumstances the same? Commercial airline pilots give each other a wide berth in the sky, reducing the need for split-second decisioning. That is certainly not the case with automobiles. And commercial airliners have two pilots, backing each other up. Not so with automobiles.

Of course, companies place great value in their branding approaches, seeking differentiation from their competitors, and seeking to attract their target market. But they also owe potential customers clarity and transparency about their products and services, knowing that some customers will not read beyond the brand name. This is all the more true when the product takes over responsibility for driving a car, thus making life and death decisions.

All three of these "pilot" systems require certain conditions to be met for the automation to take over. Tesla claims that Autopilot will steer, accelerate and brake automatically within its lane (Tesla, 2022). Presumably, anything else requires driver intervention. Mercedes Benz' initial offering of Drive Pilot was limited to specific roadways in Germany that had been mapped to exacting three-dimensional detail (Mercedes-Benz Group, 2022). Honda's initial deployment of Traffic Jam Pilot is limited geographically as well (Honda European Media Newsroom, 2021).

But how might these and other companies be held accountable for truthfully stating and delivering on their marketing promises?

In January of 2020, the U.S. Department of Transportation and the National Science and Technology Council published the fourth in a series of reports on autonomous vehicles. One of the focus areas was protecting users and communities. Emphasizing safety, the report stated that the federal government will "enforce existing laws to ensure entities do not make deceptive claims or mislead the public about the performance capabilities and limitations of AV technologies including, for example, deceptive claims relating to vehicle safety or performance" (National Science & Technology Council & U.S. Department of Transportation, 2020). There are no new federal laws yet, but there is some existing consumer protection via the Federal Trade Commission's Truth-In-Advertising statutes.

U.S. states have actively been promoting and supporting autonomous vehicle development for over a decade. As of this writing, 38 states and the District of Columbia have enacted legislation or issued executive orders regarding autonomous vehicles. The legislative actions include authorization of studies, testing, and full deployment of autonomous vehicles (Governors Highway Safety Association, 2022). Yet there is still no specific regulation regarding unsubstantiated claims from manufacturers.

Ultimately, it will be incumbent upon lawmakers and regulatory bodies to appropriately protect the public as a whole.

Questions

- Consider the range of names chosen by automakers for their self-driving technologies: e.g., GM Supercruise, Waymo Driver, Tesla Autopilot and Full Self Driving, Audi City Assist, Subaru EyeSight, Mercedes Drive Pilot. Why do automakers (and other sellers) choose names for these features that could be misleading?
- How might the automakers change their choices to improve the clarity of what the feature actually does?
- Is Melissa legally or ethically required to read and understand the feature capabilities before putting them to use?
- Thinking from one of the ethical lenses mentioned above, what are the ethical implications of the names given to these features?

Vehicle Sharing

Joseph has gone home for the night in his Level 5 autonomous vehicle. It is 7 p.m., he will have dinner, talk with the family, and turn in for the night. Meantime, his car (a $40,000+ purchase) might sit idly in the driveway all night – eight to twelve hours of underutilizing that asset. Instead, Joseph puts that car to work. He signs up with a ride-sharing service, and his car can then go out on its own and pick up riders, take them to their *destinations, and return home in time to recharge and take Joseph to his next destination.*

Since at least as far back as 1995, people around the world have started to embrace the modern version of the "sharing economy," as enabled by information technology (Puschmann & Alt, 2016). A growing sector of the sharing economy (e.g., Airbnb, Uber, Lyft, etc.) allows people to use a personal asset like a home or car to earn extra income. It may be that this same notion could extend into the use of the idle time of autonomous vehicles. Imagine you own a fully autonomous (SAE Level 5) vehicle. Beyond the nighttime rides in the story above, the car could also take the owner to work, and then it would not have to sit idle all day in the parking lot. Instead, it can pick up riders during the workday as well.

Owners might be able to choose limitations on how far the vehicle is allowed to travel, or how many people they are willing to let the car transport. For each ride that the vehicle performs, the owner earns a share of the income, while the service keeps the remainder to cover the cost of coordinating and managing the rides, riders, and vehicles. Presumably, some makers of autonomous vehicles will do better at providing highly specific navigation (getting to or from a specific building at a commercial site, etc.).

In a world where we begin sharing unsupervised access to our "personal spaces" (e.g., home, car, etc., that we routinely also use ourselves), this becomes fundamentally different from sharing space in your home (a la Airbnb) or in your car with you driving (a la Uber). The natural inclination to behave "properly" when the "renter" is with the "owner" diminishes if the owner is not present or visible, which would be the case for your car picking up rides while you sleep or work.

Questions

- What privacy issues are raised when you allow someone to use your car without your presence?
- What risks does the car owner take on in offering their vehicle this way?
- How can different car navigation capabilities be communicated to the passengers?
- What are the ethical considerations of vehicle sharing, both for the owner and for the "rider"?
- What are the pros and cons for society of enabling this sort of vehicle sharing?
- Thinking perhaps more maliciously, what other (perhaps criminal) uses might autonomous vehicles be put to?

Real Accidents with Self-Driving Vehicles

Late one night in Tempe, AZ, a test vehicle for Uber's autonomous vehicle program was traveling in fully autonomous mode with a backup vehicle operator behind the wheel. The vehicle's systems detected "something" ahead of it, moving toward the vehicle's path. The operator should have had full attention on the road, but was allegedly distracted by something else in the vehicle (Smiley, 2022). The on-board systems interpreted the "something" at various times as a bicycle, a person, and an unknown object, but did not alert the vehicle safety operator of the potential issue until it was too late, and what turned out to be a pedestrian was struck and killed.

Sadly, there is no shortage of incidents of vehicles operating autonomously (whether at Level 3 and above, fully autonomously as in this Uber case, or at Level 2 or below, where a safety driver is expected to be instantly available) that encounter accidents of some kind. The Uber incident above is just one of those. There are also numerous incidents of Tesla vehicles, running at Level 2, hitting parked emergency vehicles, perhaps due to the human-driven vehicle in front of them going around the emergency vehicle too quickly. There are also many cases where drivers have somehow bypassed the driver detection systems and let the car drive itself while the driver slept, or climbed into a different seat to "prove" how safe the car was (Dunn et al., 2021; Mak, 2021; Song et al., 2021). Other reports indicate that most accidents involving self-driving vehicles are caused by human error, such as rear-ending the autonomous vehicle (Jeffs, 2022).

Putting autonomous vehicles on the road requires the rollout of new technology that needs to operate safely on public streets, interacting with human drivers in other vehicles. Manufacturers of vehicles, and the related autonomous technology and sensors, need mechanisms and processes to safely gain confidence- and quality-building test miles without putting anyone at unreasonable risk. Even if the system is running in fully autonomous mode (e.g., Level 5), manufacturers need to test these vehicles in a public, "real-world" environment to be confident that they can trust the vehicle to make the right decisions.

Questions

- How does the operation of a vehicle at Level 4 or 5 change the dynamics of

interaction among vehicles? In other words, without a driver in control, what changes about the vehicle interactions?
- Thinking about the interactions that pedestrians have with vehicle drivers - what changes about those interactions when there is no "driver"?
- What steps could be taken by manufacturers to help mitigate the risk of putting these vehicles into the hands of the general public?

Insurance Industry Impact

In 2010, motor vehicle crashes in the United States caused 33,000 deaths, 3.9 million injuries, and 24 million damaged vehicles. The economic cost was calculated at $242 billion (Blincoe et al., 2015), including health care, lost work, repair costs, etc., and not all covered by insurance. Unfortunately, many studies have also shown that over 90% of car accidents are caused by human error (Treat et al., 1979).

One of the key anticipated benefits of ever-increasing use of autonomous vehicles is a reduction in human drivers and a presumed, corresponding reduction in accidents. That would be good, leading to a reduction in deaths, injuries, property, and financial losses.

But what happens to the automobile insurance industry? How should it change? How will it change?

Consider that the market value of the U.S. automobile insurance industry is $316 billion in 2022 (IBISWorld, 2022). That value comes from revenue, which starts with premiums charged to individual customers and businesses. Over time, with a reduction in accidents, insurance underwriters should adjust premiums downward to accurately reflect risk, thus reducing revenue.

Given that the transition to vehicles with automated driving capabilities (SAE Levels 4 and 5) will occur over many years, it is reasonable to assume that the automobile insurance industry can continue to survive and thrive while adapting to these new technologies. Discussions within the industry point to various other models for ensuring autonomous vehicles that may be explored. For instance:

No-Fault Insurance

While the track record of no-fault auto insurance is a bit bumpy, there are still 12 U.S. states that use a limited version in which drivers must use their own insurance to pay for their injuries after a crash - no matter who is at fault. They protect themselves by purchasing Personal Injury Protection (PIP) coverage (Hurst, 2022). Given that fully automated vehicles will not be driven by humans, there may be an opportunity to explore this approach.

Manufacturers May Self-Insure

Given that liability for accidents caused by autonomous vehicles will often reside with manufacturers, they may choose to self-insure - the company sets aside a pool of money to cover any losses - perhaps to a certain amount - then backed by insurance.

Fleet Insurance

It is very likely that manufacturers will partner with other businesses to manage pools (fleets) of autonomous vehicles. This business model is well known in the auto industry with, for instance, bulk sales to rental car companies and large corporations. These business relationships often leverage fleet insurance policies, which provide liability insurance for a fleet of vehicles under one policy (Stanley et al., 2020).

Questions

- What other approaches might be taken to pay for accidents involving autonomous vehicles?
- How might technological innovation aid in determining liability with autonomous vehicle accidents?
- How might the cost of insurance coverage be affected by road conditions where an autonomous vehicle is used?

3. TOPICS FOR ADDITIONAL RESEARCH

The following questions are intended to be used for additional research topics, perhaps for in-class discussion, for student research papers, or other purposes. The questions sometimes extend beyond the specifics of autonomous vehicles and into broader areas of technology deployments.

Driver (Operator) Attention

In vehicles at Level 3 automation and below, the vehicle expects the driver to be ready to take over control at a moment's notice - either because the operator detects a problem that may not have been caught by the vehicle, or because the vehicle cannot determine the correct next step. Many accidents have resulted from operators becoming too complacent and not paying attention.

Figure 3: Clinical trials (MD Anderson, 2022), used with permission

Question

- How can the automated systems in the vehicle monitor and track operator attention to ensure that the operator remains ready to take over?

Unintended Consequences of Widespread Adoption of Autonomous Vehicles

"The world is full of tradeoffs" -- Anonymous

Proponents of autonomous vehicles expect them to:

- become widely adopted,
- be less accident-prone than human drivers, and
- offer the possibility that people can readily get from place to place in an autonomous taxi, rather than owning a vehicle of their own.

If the cost of a "ride" (e.g., from an Uber-like service without a driver) becomes low enough, you will not need to own a car to be able to have car-like freedom of movement. If your own car can navigate and move entirely on its own, it may not need to park near where it drops you off - it could drive somewhere else. Perhaps it can park more cheaply elsewhere; perhaps it will pick up other passengers until you ask it to return.

Questions

- What kinds of side effects might occur if the proponents' expectations come to pass?
- What industries might see increased revenues?
- What industries might see decreased revenues?
- What new industries might appear?
- What new public safety risks does this introduce?

Parallels to Medical Clinical Trials

In some ways, the rollout of an innovative technology might be like the rollout of new medications or vaccines. Autonomous vehicles have numerous interactions with broad groups of people. Those interactions could work well in some situations and cultures, and poorly in others, with perhaps serious consequences, much like medical treatments.

Consider the process for conducting "clinical trials." In general, new treatments and approaches go through a multi-phase trial process, with carefully selected groups of participants to ensure that the tests assess the right questions, and that appropriate levels of risks are undertaken and managed.

A simple summary of the clinical trial phases follows (illustrated in Figure 3):

- Phase 1 - Tests if a new treatment is safe, and how best to deliver that treatment.
- Phase 2 - Tests if the treatment generates the desired reaction from the target. Does it generate side effects?
- Phase 3 - Tests new treatment as

compared to existing treatment (is the new treatment measurably better?).
- Phase 4 - Start broad delivery and monitor results.

For comparison, you might consider the model proposed for rollout of fully autonomous (driverless) taxi service in San Francisco, CA, run by GM's Cruise technology unit (LaReau, 2022). Their permit allows for a phased rollout of the technology, initially limited to "a maximum speed of 30 mph, from the hours of 10 p.m. to 6 a.m. daily 'when weather conditions do not include heavy rain, heavy fog, heavy smoke, hail, sleet, or snow.'" The permit could be broadened later, perhaps extending service hours, geography, or conditions. All of this is being done after the technology has been tested with human operators on board to monitor behavior and take over if needed.

Questions
- How does the rollout of autonomous vehicles compare to the clinical trial process and phases?
- Is this a potential model that could be used to structure our thinking about technology deployments to the general public?
- Are there ethical considerations for autonomous vehicle testing similar to those of clinical trials?
- Where, if anywhere, does the analogy to clinical trials break down or not make sense?
- What are the societal costs (if any) to be borne if government delays or neglects the oversight of managing the liabilities of autonomous vehicle rollouts?

4. CONCLUSIONS

Innovative technologies are in development all the time, in many sectors of the economy. Some are isolated – for example, deployment of robotic devices in a factory, which directly affects only the workers around those robots, and those whose jobs might be affected. Others, though, have a multi-faceted collection of interactions with broader groups of people, up to the broadly defined "general public." In the case of autonomous vehicles, people can be, at various times:
- owners and operators of autonomous vehicles,
- paying passengers in autonomous taxis or buses,
- users who receive shipments or deliveries from autonomous vehicles,
- driving other less-autonomous vehicles that must interact with the autonomous ones, and
- pedestrians, cyclists, and others who also must interact with autonomous vehicles.

Getting all of those interactions "right" is not a trivial problem. It is possible to design and evaluate approaches to those interactions in an isolated "laboratory" setting. Before a full and broad rollout of the technology, manufacturers (and regulators, insurers, etc.) want to be confident that those interactions will work well, and under all circumstances. Can the vehicle navigate a snowy road? Can the vehicle maneuver through a police roadblock? Does the vehicle interaction at a 4-way stop intersection work correctly? What about roundabouts? Construction sites? Detours? Do driver differences across states, countries, and cultures all work "correctly"?

Gaining enough confidence that those approaches are correct requires many hours of interactions with large swaths of the population, to be sure that the interactions work, without fail, under all conditions. Those hours require taking a variety of risks, and it is critical for technology managers to think through those risks and engage the other participants in the process to mitigate those risks appropriately. We hope that these scenarios have provided things to consider about the challenges of delivering new technologies to the public, and about managing those risks appropriately.

Governments and regulatory agencies also need to define the legal framework in which these vehicles will be operated, during testing and after full rollout. Legal liabilities should be well-defined and appropriate for all the potential scenarios. Regulators should carefully consider the full range of risk tradeoffs and specify appropriate compensating controls to offset each of those risks. Those controls might include limiting the conditions in which these systems can be used, limiting the geography where they can be used, or other similar constraints to manage the risk to the public. Automakers need to be crystal-clear about, and actively enforce, the need for driver attention. Only then can we safely get from "I drive" to "it drives."

5. ACKNOWLEDGEMENTS

The authors appreciate the support and insights of various industry experts who provided insights for this case. The authors also appreciate the productive feedback provided by their peers, and by the reviewers and conference chairs.

6. REFERENCES

Ahmed, R. (2017). Risk Mitigation Strategies in Innovative Projects. In *Key Issues for Management of Innovative Projects*. IntechOpen. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3884687

Azuma, M. C., Giordano, F., Stoffregen, S., Klos, L., & Lee, J. (2023). It practically drives itself: autonomous vehicle technology, psychological attitudes, and susceptibility to risky driving behaviors. *Ergonomics 66*(2), 246-260.

Blincoe, L. J., Miller, T. R., Zaloshnja, E., & Lawrence, B. A. (2015). *The Economic and Societal Impact Of Motor Vehicle Crashes*. U.S. Department of Transportation, National Highway Traffic Safety Administration https://crashstats.nhtsa.dot.gov/Api/Public/ViewPublication/812013

CBS News Bay Area. (2021). *Dead-End SF Street Plagued With Confused Waymo Cars Trying To Turn Around 'Every 5 Minutes'*. Retrieved June 9 from https://www.cbsnews.com/sanfrancisco/news/dead-end-sf-street-plagued-with-confused-waymo-cars-trying-to-turn-around-every-5-minutes/

Dunn, N. J., Dingus, T. A., Soccolich, S., & Horrey, W. J. (2021). Investigating the impact of driving automation systems on distracted driving behaviors. *Accident Analysis & Prevention*, *156*, 106152.

Goodall, N. J. (2021). Comparison of automated vehicle struck-from-behind crash rates with national rates using naturalistic data. *Accident Analysis & Prevention*, *154*, 106056. https://profile.virginia.edu/sites/g/files/jsddwu351/files/njg2q/struckfrombehind_1.pdf

Governors Highway Safety Association. (2022). *Autonomous Vehicles*. Retrieved June 8 from https://www.ghsa.org/issues/autonomous-vehicles

Greenberg, A. (2016, 1 Aug). The Jeep Hackers Are Back to Prove Car Hacking Can Get Much Worse. *WIRED*. https://www.wired.com/2016/08/jeep-hackers-return-high-speed-steering-acceleration-hacks/

Honda European Media Newsroom. (2021). *Honda launches next generation Honda SENSING Elite safety system with Level 3 automated driving features*. Retrieved May 27, 2022, from https://hondanews.eu/eu/et/corporate/media/pressreleases/329456

Hurst, A. (2022). *States with no-fault auto insurance and what it means*. PolicyGenius. Retrieved May 30, 2022, from https://www.policygenius.com/auto-insurance/states-with-no-fault-insurance/

IBISWorld. (2022). *Automobile Insurance in the US - Market Size 2004–2027*. Retrieved May 25, 2022, from https://www.ibisworld.com/industry-statistics/market-size/automobile-insurance-united-states/

Jeffs, D. J. (2022). *Autonomous Cars, Robotaxis & Sensors 2022-2042*. IDTechEx. https://www.idtechex.com/en/research-report/autonomous-cars-robotaxis-and-sensors-2022-2042/832

Lambert, F. (2021, May 5). Tesla driver repeatedly spotted in backseat on Autopilot is begging to be arrested. *Elektrek*. https://electrek.co/2021/05/05/tesla-driver-keeps-being-spotted-in-backseat-autopilot-begging-arrested/

LaReau, J. L. (2022, June 2). GM's Cruise will be the first to offer driverless taxis in a major U.S. city. *Detroit Free Press*. https://www.freep.com/story/money/cars/general-motors/2022/06/02/gm-cruise-approved-california-driverless-taxis/7487221001/

Litman, T. (2017). *Autonomous vehicle implementation predictions*. Victoria Transport Policy Institute Victoria, BC, Canada.

Mak, A. (2021, Aug 17). Why Teslas Keep Striking Parked Firetrucks and Police Cars. *Forbes*. https://slate.com/technology/2021/08/teslas-allegedly-hitting-emergency-vehicles-why-it-could-be-happening.html

Malan, D. (2018). The law can't keep up with new tech. Here's how to close the gap. https://www.weforum.org/agenda/2018/06/law-too-slow-for-new-tech-how-keep-up/

Markkula Center for Applied Ethics. (2021). *A Framework for Ethical Decision Making*. Santa Clara University. Retrieved June 10, 2022, from https://www.scu.edu/ethics/ethics-resources/a-framework-for-ethical-decision-making/

MD Anderson Cancer Center. (2022). *Phases of Clinical Trials*. Retrieved June 4, 2022, from https://www.mdanderson.org/patients-family/diagnosis-treatment/clinical-trials/phases-of-clinical-trials.html

Mercedes-Benz Group. (2022). *The front runner in automated driving and safety technologies*. Retrieved May 30, 2022, from https://group.mercedes-benz.com/innovation/case/autonomous/drive-pilot-2.html

metamorworks. (2022). Autonomous Car Remote Sensing System. In 530214349 (Ed.). Shutterstock.com.

National Science & Technology Council, & U.S. Department of Transportation. (2020). *Ensuring American Leadership in Automated Vehicle Technologies - Automated Vehicles 4.0*. The Office of Science and Technology Policy. https://www.transportation.gov/av/4

Puschmann, T., & Alt, R. (2016). Sharing economy. *Business & Information Systems Engineering*, *58*(1), 93-99.

Smiley, L. (2022, April). "I'm the Operator": Aftermant of a Self-Driving Tragedy. *WIRED*, *30*(04), 54-69. https://www.wired.com/story/uber-self-driving-car-fatal-crash/

Society of Automotive Engineers International. (2021, May 23). SAE Levels of Driving Automation™ Refined for Clarity and International Audience. *SAE Blog*. https://www.sae.org/blog/sae-j3016-update

Song, Y., Chitturi, M. V., & Noyce, D. A. (2021). Automated vehicle crash sequences: Patterns and potential uses in safety testing. *Accident Analysis & Prevention*, *153*, 106017.

Stanley, K. D., Grise, M., & Anderson, J. M. (2020). *Autonomous Vehicles and the Future of Auto Insurance*. R. Corporation. https://www.rand.org/pubs/research_reports/RRA878-1.html

Szymkowski, S. (2020). Audi hangs up hopes for Level 3 partial automation system. Retrieved May 31, 2022, from https://www.cnet.com/roadshow/news/audi-a8-level-3-automation-traffic-jam-pilot-system/

Tesla. (2022). *Autopilot*. Retrieved May 27, 2022, from https://www.tesla.com/autopilot

Treat, J. R., Tumbas, N. S., McDonald, S. T., Shinar, D., Hume, R. D., Mayer, R. E., Stansifer, R. L., & Castellan, N. J. (1979). *Tri-Level Study of the Causes of Traffic Accidents*. https://deepblue.lib.umich.edu/handle/2027.42/64993

U.S. Department of Transportation. (2021). *Automated Vehicles Comprehensive Plan*. https://www.transportation.gov/sites/dot.gov/files/2021-01/USDOT_AVCP.pdf

Wendt, Z., & Cook, J. S. (2018). *Saved by the Sensor: Vehicle Awareness in the Self-Driving Age*. Retrieved June 9, 2022, from https://www.machinedesign.com/mechanical-motion-systems/article/21836344/saved-by-the-sensor-vehicle-awareness-in-the-selfdriving-age

Editor's Note:

This case was selected for inclusion in the journal as the EDSIGCON 2022 Best Case The acceptance rate is typically 2% for this category of case based on blind reviews from six or more peers including three or more former best papers authors who did not submit a case in 2022.