Autonomous Vehicles for Personal Transport: A Technology Assessment

This article was written by Matthew Moore and Beverly Lu in partial fulfillment for a course entitled "Management of Technology" taught by Ken Pickar. Both Matthew and Beverly are PhD candidates at Caltech in chemistry. They are both very interested in the processes by with promising innovative technologies cross into commercialization.

I Was Promised Flying Cars

Whether it was flying cars or jetpacks we were all promised a compelling vision of the future at some point in our childhoods. And, although it may seem like science fiction, autonomous vehicles—cars that drive themselves—are more than just a promise; they are a growing reality. The widespread coverage in October 2010 of Google's fleet of Toyota Prii that continue to navigate the highways of California was a clear message to the world that this technology is fact not fiction[1-3]. Questions remain however as to whether or not the world and the marketplace are ready for this technology.

Google's 2010 shenanigans cannot help but invoke a memory of the 1939 World's Fair as it was vividly described in Randal O'Toole's 2009 book, *Gridlock: Why We're Stuck in Traffic and What to Do About It*[4]. There is a chapter in this book—aptly named "Dude where's my driverless car?"—that depicts GM's forward-looking exhibit from the 1939 World's Fair on autonomous highways as a transparent message to the public. Namely, if you build a system of interstate highways this *could* be your future—and we did.

Today, in a very similar way, Google is saying to the world, "This *is* the future, and we will help you reach it." In fact, their feelings on the matter are clearly evidenced by Google's own recent hiring of a lobbyist in Nevada to help pass legislation that would allow automatous vehicles to roam the roadways of that state[5,6]. Still, although Google's intentions are quite transparent, an unavoidable issue at hand is to assess where this technology fits into today's personal vehicle marketplace.

To explore the topics at hand, we conducted an assessment of the current status of autonomous vehicle technologies in the U.S to the best of our own extent. We focused mainly on those technologies that are applicable to modern cars driven on existing, unmodified public roads. Our geographical focus was mainly on the west coast of the United States where a hotbed of activity regarding this technology is located.

That State of the Autonomous Art

A large portion of the recent progress in autonomous vehicle technology can be directly attributed to two competitions staged by the Defense Advanced Research Projects Agency (DARPA). In one of these programs, the 2005 DARPA Grand Challenge, competitors were tasked with building a driverless vehicle that could navigate rough, unpaved desert terrain similar to a battlefield; a group from Stanford won this event[7]. In

the other program, the 2007 DARPA Urban Challenge, teams were confronted with a full-scale mock urban environment complete with other human drivers on the course. Professor Christopher Urmson's team from Carnegie Mellon—in collaboration with General Motors (GM)—took home the grand prize[8].

Technical Opinions from the Experts

It is fair to say that there are still significant limits to the state-of-the-art in autonomous driving. To assess these issues from an academic perspective, we interviewed David Stavens, a PhD scientist who obtained his doctorate under Sebastian Thurn—the director of Stanford's Artificial Intelligence Laboratory. In the course of his PhD research, Stavens helped put together a driverless Volkswagen Touareg wagon that won the 2005 DARPA Grand Challenge.

From Stavens point of view, the biggest problems in autonomous vehicle technology lie in sensor perception and decision-making under conditions of uncertainty. These situations can include instances when sunlight shines directly into a camera or in the presence of heavy fog and mist. On the other hand, Stavens was quick to point out that no major technical hurdles still need to be surmounted for this technology to be implemented in cars today.

In addition to Stavens we also interviewed Richard Murray, a professor of control & dynamics systems at Caltech, and we met with a postdoctoral scholar in Murray's lab who led the team responsible for Caltech's entry in the 2007 DARPA Urban Challenge. Murray noted that autonomous vehicles are a "doable" technology and that the DARPA challenges were extremely useful for demonstrating that research done over the past 20 years into this kind of technology makes autonomous vehicles possible.

Murray's postdoc, Noel du Toit, explained that specific isolable problems solved in the course of research and engineering for the DARPA challenges have turned into the "driver assistance" technologies that we see in cars today. These technologies include things like automatic parallel parking, blind-spot vehicle detection, adaptive cruise control, and emergency braking for collision avoidance[9].

Still, du Toit reiterated that there are still problems with autonomous vehicle technology and "solving all the problems all the time is still a challenge." Specifically, du Toit stated that there are issues lying in the robustness of the technology since very minor changes in a vehicle's environment can throw off its entire perceptive system.

To get a commercial R&D perspective on the issue, we communicated with Dmitri Dolgov, a software engineer for Google working on autonomous driving. Dolgov was pictured behind the wheel of one of Google's retrofitted Toyota Prii in John Markoff's late 2010 article for the New York Times about Google's fleet of autonomous vehicles[1]. Dolgov's words echoed the problem that Stavens discussed with obtaining accurate and reliable perception under all conditions: nighttime, daytime, rain, sleet, and snow.

Comparisons to Other Transportation Technologies

For purposes of comparison to other types of transportation technologies, we chose a variation of the driver-miles-driven per fatal event metric used by the U.S. bureau of transportation statistics[10]. In the case of autonomous vehicles, the relevant incident to measure cannot be "fatal events" since driverless cars are not a currently available commercial product, and consequently, there has been no opportunity to collect such data. Instead, we used a metric suggested by David Stavens known as the "mean failure distance". This distance is equivalent to the average number of autonomous miles driven per required human intervention.

The term "required human intervention" as it appears in here is not meant to convey a situation that would necessarily result in a fatality. For example, Stavens described a situation where, upon going around a sharp turn in the roadway, an autonomous vehicle would lose track of the car in front of it since the preceding car exits the perceptive field of the autonomous vehicle. Under a circumstance like this, the autonomous vehicle may speed up in response to an alleviation of its following distance requirements. Such a behavior would require a human intervention to correct since it is known that the preceding car did not, in fact, disappear; it is just temporarily out of view.

Regardless of the event that results in a "required human intervention", we can say that in the best cases the vehicle will be able to correct its behavior in time to prevent a collision. In the some cases though, an unaddressed need for "required human intervention" may lead to something graver. However, we feel that we can use this metric as a stringent comparison against known fatality-related data for various travel systems.

For illustrative purposes, we have chosen to depict the status of autonomous driving technology on an industry-wide technology S-curve (see **Figure 1**). At present the mean failure distance of autonomous vehicles is on the order of 1000s of miles. Press statements made by Google concerning its efforts with the Toyota Prius, as well as by reports from VisLab in Europe on their autonomous van trek from Italy to Shanghai for the 2010 World Expo corroborate this figure[2,11,12].

Stavens suggested that at a mean failure distance on the order of millions of miles, the technology would be in a more commercially viable position, leaving it poised for mainstream use. We have chosen this distance as an inflection point for the technology since, as a technological target, it is consistent with the "six sigma" concept of quality (i.e. 3.4 failures per million). This number can also be compared with the number of driving miles per fatal event for cars in 2008, which is on the order of 100 million miles[10]. This comparison assumes that one in every hundred "required human interventions" would lead to a fatality, and the validity of this assumption is unknown.

Figure 1: Technological S-curve for mean failure distance (driving miles per required human intervention) for autonomous vehicle technology.

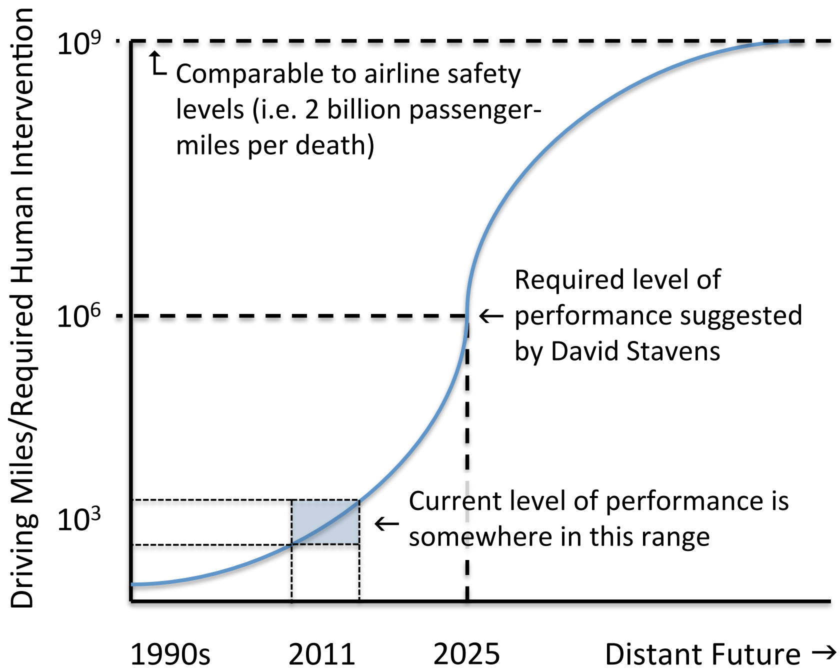

In terms of predictions, we chose to assume that the technology empowering autonomous vehicles could improve at least as fast as microprocessors. This assumption is based on the idea that autonomous vehicle technology has been largely powered by the advent microprocessor technology. Given that assumption, we used Moore's law to shape our estimates. Thus, supposing that the relevant measure is the average driving distance between required human interventions, and if that distance were to double every 18 months from what was reported at the end of 2010, it would take 15 years to get to a 1.0 million mile target. That would put the inflection point of our S-curve at about 2025.

For an estimate of the terminus for our technology S-curve, we chose a mean failure distance of billions of miles. This number is comparable to the oft-quoted airline industry safety statistic of 2.0 billion passenger-miles per fatal event. We reasoned that when driving safety in autonomous vehicles is comparable to the safety levels of air travel, the marginal cost of improving autonomous vehicle technology in terms of mean failure distance would exceed any potential marginal benefit.

A Mixture of Fascination and Skepticism

In our interview with him, Congressman Adam Schiff described the public's perception of autonomous driving technologies as a reflection of his own reaction to the idea: one that is a mixture of both fascination and skepticism. Schiff explained that the public's fascination comes from amazement at how advanced this technology already has become, plus with Google's sponsorship and endorsement it becomes even more alluring.

This fascination is somewhat to be expected given the tight alignment of autonomous vehicle technology with our collective fantasy of the future. Popular culture examples of autonomous vehicles abound, and they include KITT from Knight Rider and Herbie the Love Bug, which even border on anthropomorphic. Thus, seeing driverless cars in action is akin to looking at a vision of our own imagined future, or at least a fantasy come to life.

Skepticism of autonomous vehicle technologies comes from a missing element of trust. According to Clifford Nass, a professor of communications and sociology at Stanford University, this trust is an aspect of public opinion that must be earned through demonstration more so than through use. When people see a technology in action, they will begin to trust it.

Professor Nass specializes in studying the way in which human beings relate to technology, and he has published several books on the topic including *The Man Who Lied to His Laptop: What Machines Teach Us About Human Relationships*. In our interview with him, Professor Nass explained that societal comfort with technology is gained through experience, and acceptance occurs when people have seen a technology work enough times collectively. He also pointed out that it took a long time for people to develop trust in air transportation, something that we almost take for granted now.

It is certainly not the case that autonomous cars need to be equivalent in safety to plane flight before the public would adopt them. However, as Noel du Toit pointed out, we have a higher expectation for autonomous cars than we do for ourselves. Simply put, if we are willing to relinquish the "control" over our vehicles to an autonomous power, it will likely have to be under the condition that the technology drives more adeptly than we ever possibly could. Otherwise, there will simply be no trusting it.

Interestingly, du Toit brought up a recent botched safety demonstration by Volvo in May of 2010. In the demonstration, Volvo showcased to the press how its emergency braking system works as part of an "adaptive cruise control" system. These systems allow a driver to set both a top speed and a following distance, which the vehicle then automatically maintains. As a consequence, if the preceding vehicle stops short, the system acts as the foundation for an emergency-braking maneuver. However, In Volvo's demonstration the car smashed directly into a trailer[13]. Even though the system worked fine in several cases during the day's worth of demonstrations, video of that one mishap went viral and did little to help the public gain trust in the technology.

We asked Professor Nass to compare public perception of autonomous driving to an autonomous vacuuming technology such as the Roomba®. In his response, Professor Naas clarified that the Roomba® does essentially exactly what we would expect it to do in the way we expect it to do so; the notion of the technique is very straightforward. This however is not true of driving. Nass went on to say that since humans are scared of technology that they do not understand, conveying the logic of why an autonomous vehicle is making certain decisions would be very important.

Besides, it is one thing to program in an arbitrary set of rules into the bug-like intelligence of a robot like the Roomba®. However, it is entirely another thing to encode a set of canonical guidelines for a vehicle that is expected to pilot the roadways with other human drivers. David Stavens found this to be the case when he noted that an attempt at hard-coding these rules just ended up meaning adding exception upon exception until the whole web of rules was too complicated for even a computer to navigate efficiently.

What Stavens found in the course of his research was that allowing the vehicles to write their own code through machine learning was much more effective. Essentially, a vehicle would "observe" human drivers enough to come up with it's own self-consistent set of instructions. It turns out this approach was much more efficient and even made more logical sense than one might have expected from collective human faith in the "rules of the road".

Legalities and Liabilities (i.e. Speed-bumps)

To get a better sense of the legal issues surrounding autonomous vehicle technology, we interviewed Ryan Calo, the director of the Consumer Privacy Project at Stanford Law School's Center for Internet and Society. Calo made it clear that the legal issues behind autonomous driving are not fully sorted yet, but explained that the legal code relating to this technology essentially boils down to the following: "every vehicle needs to have a licensed operator."

Calo pointed out at that future issues related to autonomous vehicles would be approached from a standpoint of "negative liabilities", meaning that we can assume something is legal unless there exist explicit laws against it. This discussion also led to the concept of what a driverless car would look like to bystanders, and the kind of panic that might garner. A real-life example of this occurred in Moscow during the VisLab van trek to Shanghai[11]. In this case, an autonomous electric van was stopped by Russian authorities due to its apparent lack of a driver behind the wheel. Thankfully, engineers present were able to convince the Russian officer who stopped the vehicle not to issue a ticket.

A current legal ambiguity that Calo pointed out refers to the status of technologies like adaptive cruise control and lane-keeping assistance. The latter is a technology based on cameras that monitor painted driving lanes. The end result being that the car provides some resistance to turning the steering wheel if it begins to change lanes without signaling. Combining adaptive cruise control with lane-keeping assistance results in a pseudo-autonomous driving situation, and one could potentially argue, a perfectly safe environment in which to begin using one's cellular phone.

In regards to Google's recent lobbyist activity in Nevada, Calo clarified that the bill under consideration in that state aims to shift licensing from the driver to the vehicle[14]. This would provide a work around to the stipulation that "every vehicle needs to have a licensed operator" because the vehicle would be its own operator. Accordingly, as Calo

explained, a set of driving tests would be implemented for autonomous vehicles in a fashion similar to those used to test human drivers. Calo was uncertain as to whether or not follow-up testing would be required after license expiration.

The above legislation fits in well with the information that we collected from Congressman Schiff about potential federal involvement in autonomous vehicle technology. Basically, Schiff relayed the idea that a strong governmental role expected for this technology would come in the form of regulating safety. Furthermore, he called attention to hefty governmental requirements for crash testing that every new vehicle must meet before it is allowed on the road.

With respect to the Nevada bill, Calo made one thing clear: the bill removes the concern of “lawfulness” related to autonomous vehicles, but it does not purport to do anything about issues related to liability. Liabilities are nothing new to both the auto industry and the robotics industry. Professor Murray at Caltech cited large liability policies taken out by automakers to prepare for the introduction of airbags into vehicles. He also mentioned that use of surgical robots was a place in the field of robotics where liability concerns had been encountered in the past.

In autonomous driving, liability concerns can be inferred through a couple of examples. In one example, Noel du Toit described DARPA’s use of hired stunt drivers to share the testing grounds with driverless vehicle entries in the 2007 Urban Challenge. This behavior clearly illustrates the level of precaution that the DARPA officials felt it was necessary to take. In another example, Dmitri Dolgov expounded on how Google’s cars are never driving by themselves; whenever they are operated on public roads, there are at least two well-trained operators in the car. Dolgov went on to say that these operators “are in control at all times”, which helps illustrate Google’s position—they are not taking any chances when it comes to liabilities.

Kent Kresa, former CEO of Northup Grumman and interim chairman of GM in 2009, was also concerned about the liability issues presented by autonomous vehicles. Kresa felt that a future with driverless cars piloting the streets was somewhat unimaginable at present, especially when one considers the possibility of a pedestrian getting hit. In the case of such a collision it is still very unclear who would be at fault. Whether or not the company that made the vehicle would be responsible is at present unknown.

The Road to Adoption (Pun Intended)

There are several options for the imagined path to commercialization of autonomous vehicles. The technical experts we interviewed were primarily in consensus that the technology for fully autonomous vehicles would be implemented gradually; these people included David Stavens, Dmitri Dolgov, and Noel du Toit. Congressman Schiff was also able to envision this path as a possible future for autonomous vehicles.

In this gradual adoption model, technologies advertised under the guise of “driver assistance” or “advanced safety features” would gradually supplant the normal operations

of driving a vehicle. In **Table 1**, we provide a sample of options packages available on current vehicles that include some of these features. A cursory viewing of the prices for these packages suggests that car-buyers are willing to pay for them with a premium of 5–10% the MSRP of their vehicle. According to Stavens however, the total cost for a typical sensor suite on one of Google's autonomous Toyota Prii would be between $100 and $150 thousand. Currently, such a price point would be prohibitive.

Table 1: This table provides a small sampling of known options packages obtainable for 4-door sedan type personal vehicles available for purchase in 2011 that include one or more of the following: adaptive cruise control, lane-keeping assistance, collision warning, blind-spot monitoring, and parking guidance.

Vehicle	MSRP	Option Name	Price
Acura RL	$31,600	Advance Package*	$3,800
Audi A8	$78,050	Driver Assistance Package	$3,000
Audi A7	$47,000	Side Assist Package	$500
BMW 750i	$70,650	Driver Assistance Package	$3,500
Dodge Charger Rallye	$32,005	Adaptive Cruise Control Group	$925
		Driver Confidence Group	$1,495
Infiniti M56X	$61,600	Technology Package	$3,000
Mercedes-Benz S550	$93,000	Driver Assistance Package	$2,950
Toyota Prius 5	$23,050	Advanced Technology Package†	$5,080
Volvo S60	$30,975	Technology Package	$2,100

* Calculated from difference in price of model with "Technology Package"

† Includes price of the GPS navigation system ($2,380)

In another model, autonomous vehicles would be introduced through special dedicated lanes of traffic. Brian Thomas, managing director at the Public Financial Management group, pointed out that the construction of this type of infrastructure could be funded through "public-private partnerships (P3)." Thomas noted that toll roads in Indiana and Chicago have already been built using P3 funding. He also thought it was not infeasible for a private firm to buy all the HOV (High Occupancy Vehicle) lanes in the Los Angeles area for this purpose. However, Thomas speculated that such a move would be subject to strong political pushback due to feelings of discrimination against lower income brackets.

Noel du Toit opined that big challenges lay ahead in the future for integrating the use of autonomous cars with non-autonomous human drivers. For this reason, it would make sense initially to have the technology be switchable on-and-off. In the dedicated lane model of commercialization, it could be envisioned that it would be a requirement to switch into autonomous mode once a vehicle enters the "autonomous-only" lanes.

Significant effort was spent in the 1990s to test the dedicated lane adoption model on an extension of interstate 15 into San Diego[15]. In this case, magnets were installed at regular intervals in the HOV lane of the freeway and a platoon of eight Buick LaSabres was run along a 7.6 mile stretch of the roadway at 65 miles per hour, all of them merely 18 feet

apart each. The cars were not fully autonomous and communicated with each other through radio frequency transponders extensively.

Since then, the burden of responsibility has shifted from government-sponsored infrastructure aimed at enabling this technology through an “automated highway system” to a palette of cars that are themselves “smarter”[4]. This paradigm shift may be due to financial concerns. As Brian Thomas confirmed, city governments by themselves face large financial barriers to maintaining infrastructure, let alone updating it.

A conversation we had with Bruce Gillman, the public information officer for the Los Angeles Department of Transportation (DOT), revealed that the department is very busy putting out many other fires. Gillman noted that DOT is focused on getting people out of their cars and onto bikes or into buses. Thus, autonomous vehicles are not on their radar. Moreover, Gillman was adamant that DOT would wait until autonomous vehicles were being manufactured commercially before addressing any issues concerning them. His viewpoint certainly reinforces that idea that supportive infrastructure updates coming form a city government level would be unlikely.

No matter what adoption pathway is used, federal government financial support could come in the form of incentives and subsidies like those seen during the initial rollout of hybrid vehicles. However, Brian Thomas explained that this would only be possible if the federal government was willing to do a cost-benefit valuation for the mainstream introduction of autonomous vehicles.

Besides the gradual adoption model and the dedicated lane model, we can only speculate on other pathways to commercialization. Public transportation vehicles may seem like an obvious choice given their defined routes and predictable usage patterns[16]. In fact, Siemens recently made a big deal to provide equipment for a driverless subway system in São Paulo[17]. Furthermore, automated buses reportedly shuttle workers in the Dutch city of Rotterdam from the Kralingse Zoom metro station to a nearby office park[18]. In fact, there was one highly publicized crash between these buses when no passengers were on board, but this was later determined to be due to human error[19].

Another possible entry route for the technology is through platoons of freight vehicles like semi-trucks[20]. According to Noel du Toit, extensive work has been done on these sorts of systems in Europe since the mid-90s[21-23]. In this case, a single truck driver would pilot a vehicle that is followed closely by a chain of driverless trucks using systems similar to the adaptive cruise control and lane-keep assistance technologies described earlier. Initially, the following trucks need not be driverless, but rather different drivers could take turns as platoon leader while others rest.

One place where commercial realities for autonomous driving are apparently being explored is in services for the elderly. Professor Murray relayed that the Japanese company Hitachi approached him at some point to discuss employment of driverless vehicle technology for use in delivering groceries to the elderly. According to Murray, Hitachi had sized this potential market in Japan at approximately 400,000 vehicles.

With extensive advances in current technology, it may even be possible to enter commercialization through an autonomous taxi system for the confined downtown area of a city[9]. As Ryan Calo pointed out though, right now this idea is a bit unrealistic given the current state-of-the-art. However, as Brian Thomas opined, it would be a stellar way to get to and from the airport in many cities.

Next Stop, The Future

Visions of the Future

Many of our interviewees had dramatic and captivating visions of a future world filled with autonomous vehicles. For example, David Stavens envisioned the end of sprawling giant parking lots at university campuses and airports. Relatedly, Ryan Calo imagined a future where his car could be doing something else for eight hours a day like running errands or carrying around other passengers instead of being parked. In this vision, it is even possible to conceive of smart phones putting for-rent autonomous vehicles at one's beck and call.

One thing is for sure; a system of driverless vehicles would change to the global car culture[24]. Professor Sebastian Thurn noted in his 2011 TED Talk that he looks forward to a future where people look back and wonder why humans ever thought to drive cars themselves[3]. While that future is probably more than 50 years off, it is plausible to think that widespread use of autonomous vehicles will happen, for safety concerns alone.

David Stavens opined that sufficiently reliable autonomous vehicles could even compete with the airline industry for short flights. For example, Stavens noted that an autonomous vehicle driving all night at 50 miles per hour could travel close to 500 miles. Effectively, this could mean home-to-hotel overnight service for a given passenger. A service like this may end up being cheaper and more convenient than plane flight in a variety of circumstances.

Alternatively, Tom Zambrano, director of technical innovation at AeroVironment, conveyed to us that he thought autonomous cars bring little to the table in our future. Zambrano pointed out that the last unexplored places on Earth are deep in the ocean, and autonomous underwater vehicles represent a truly compelling future direction for autonomous vehicle technology.

Kent Kresa and Congressman Adam Schiff both imagined pilot-less aircraft to come before driverless cars. From a technological standpoint, this makes sense. At present, a large portion of airline flight is done through "auto-pilot" systems. Also, there is much more room for error in the air than on the ground; being off by a meter in the air is not the difference between the road and the sidewalk. Zambrano also pointed out that even pilot-less flying drones used by the U.S. military do not take advantage of their autonomy. After all, these vehicles only go directly from point A to point B whereas they could have the capacity to sense wind direction and ride thermal updrafts for even greater efficiency.

Professor Clifford Nass was very skeptical that the general public would be willing to fly anywhere in a pilotless plane, even if they knew it could be remote controlled from the ground. Similarly, Professor Richard Murray voiced that even a future containing a large system of semi-autonomous trains should still have conductors. Murray also noted that autonomous vehicles for personal transport seemed like a particularly unsustainable solution for the future of the human race. And, while some further fuel efficiency may get eked out of autonomous driving, Murray's point is valid.

Tom Zambrano also described to us a future vision of goods and package transport run by smaller-scale autonomous carrier vehicles. So, the existence of armies of autonomous robotic droids swarming the roads and the skies is one possible vision of the future as well[24].

Predictions

The future of autonomous vehicle technology is hazy in some respects and concrete in others. Congressman Schiff speculated that if Google is already involved in autonomous vehicles, then this technology might be closer to commercial viability than one would expect. However, what is concerning at present is that the only heavy hitting organization openly involved in this technology is apparently Google. If other large automakers were openly endorsing or developing the technology, perhaps one could say with more confidence something about its implementation in the next five to ten years.

History shows that wild projections about autonomous vehicles are not well received. Close to its impending doom in 2008, GM's former CEO Rick Wagoner predicted commercial production of an autonomous car by 2018[25,26]. In retrospect, his statement does not seem completely unsubstantiated. After all, in 2007 it was the GM collaboration with Carnegie Mellon that won the DARPA Urban Challenge.

What has occurred and is likely to continue to occur is the implementation of component technologies. Just as anti-lock braking systems, automatic headlights, and air conditioning eventually became factory standards on all personal vehicles, so will components of autonomous vehicle technology that represent market value to consumers. Cruise control is present on the base models of most modern vehicles; it is not inconceivable that adaptive cruise control could become standard in ten years. At the very least, safety benefits derived from emergency braking technology, which is a consequence of adaptive cruise control, would likely be astounding.

On the other hand, the push for these advances will have to come from the market. Right now, many of these driver assistance components of autonomous vehicles are only available as options for high-end models of cars on the market. That needs to change. Equipping a car with emergency braking and adaptive cruise control has nothing to do with that car having leather seats or gold trim. These technologies are mainly for safety, not solely for luxury or convenience, and their market adoption will only by hastened by making them more widely available.

Besides, as the consumer demand for increased usage of mobile devices increases, the value proposition of a car that drives itself will skyrocket. This would be especially true if text messaging was allowed while driving vehicles equipped with these technologies. In relation to this point, part of the Nevada bill on autonomous vehicles seeks an exemption to allow texting for riders of autonomous vehicles[5,6].

The day when driverless vehicles will roam the streets freely to pick up and drop off passengers like futuristic taxies is far off, probably more than 20 years at least. However, the day when most drivers on the highway will not really be driving themselves at all is closer than anyone may think—perhaps 10 years or even less.

References

1. Markoff, J., "Google Cars Drive Themselves, in Traffic", *The New York Times*, October 9 2010
2. Siegler, M. G., "Google Has a Secret Fleet of Automated Toyota Priuses; 140,000 Miles Logged So Far", *techcrunch.com*, October 9 2010
3. Thurn, S., "Google's Driverless Car", *TED talk*, March 2011
4. O'Toole, R., *Gridlock: Why We're Stuck in Traffic and What to do About It*, Ch. 9: "Dude, Where's My Driverless Car?", 2009, The Cato Institute.
5. Markoff, J., "Google Lobbies Nevada to Allow Self-Driving Cars", *The New York Times*, May 10 2011
6. Calo, R., "Nevada Bill Would Pave the Road to Autonomous Cars", *cyberlaw.stanford.edu*, April 27 2011
7. Mohammed, M., "'Junior' on Track for Urban Race", *The Stanford Daily*, October 26 2007, Stanford University
8. Urmson, C. *et al.*, "Autonomous Driving in Traffic: Boss and the Urban Challenge", *Artificial Intelligence Magazine*, Summer 2009, Association for the Advancement of Artificial Intelligence
9. "The Self-Driving Car Comes Closer—But Difficulties Remain", *The Economist*, August 25 2005
10. U.S. Bureau of Transportation Statistics. Available from: http://www.bts.gov
11. Kent, J. L., "Driverless Van Crosses from Europe to Asia", *CNN*, October 28 2010
12. Kurtenbach, E., "Without Driver or Map, Vans Go From Italy to China", *USA Today*, October 28 2010, Associated Press
13. Lanxon, N., "Volvo Smashes Car in Safety Demo", *Wired Magazine*, May 7 2010
14. "Assembly Bill No. 511-Nevada Committee on Transportation", March 28 2011
15. Ashley, S., "Smart Cars and Automated Highways", *Mechanical Engineering Magazine*, May 1998, The American Society of Mechanical Engineers
16. Parks, B., "Robot Buses Pull In to San Diego's Fastest Lane", *Wired Magazine*, July 24 2007
17. Fisher, D., Karmali, N. & Epstein, G., "Urban Outfitter", *Forbes Magazine*, May 9 2011
18. de Wolf, A., "Driverless Robot Buses Crash", *wolfstad.com*, December 6 2005

19. de Wolf, A., "Driverless Robot Buses Crash, Part 2", *wolfstad.com*, December 17 2005
20. Ward, M., "'Road Train' Technology Trials Get Rolling", *BBC News*, January 18 2011
21. Hooper, S., "The Car That Can Read Road Signs", *CNN Tech*, October 7 2004, CNN
22. Strong, D., "Bold New Vectra to Drive Itself!", *Auto Express News*, October 18 2005, Auto Express
23. "Self-Parking Cars Coming to U.S.", *Orlando News*, April 4 2006, Click Orlando
24. Birch, A., "Most Cars Can Be Eliminated In 20 Years, Says Urban Designer Michael E. Arth", *corrupt.org*, May 23 2008
25. Squatriglia, C., "GM Says Driverless Cars Could Be on the Road by 2018", *Wired Magazine*, January 7 2008
26. Stoll, J. D., "Could GM's Salvation Be Stuff of Science Fiction?", *The Wall Street Journal*, January 7 2008

List of Interviewees
(ordered by date of interview)

1. **David Stavens,** *PhD Scientist at Stanford Artificial Intelligence Laboratory*
2. **Kent Kresa,** *Former CEO of Northrop-Grumman, Interim Chairman of GM*
3. **Dmitri Dolgov,** *Software Engineer at Google in the Autonomous Driving*
4. **Tom Zambrano,** *Director of Technical Innovation at AeroVironment*
5. **Clifford Nass,** *Professor of Communications and Sociology at Stanford*
6. **Bruce Gillman,** *Public Information Officer at LA Dept. or Transportation*
7. **Adam Schiff,** *U.S. Congressman for California's 29th District*
8. **Richard Murray,** *Professor of Control & Dynamical Systems at Caltech*
9. **Noel du Toit,** *Post-Doc in Control & Dynamical Systems at Caltech*
10. **Ryan Calo,** *Director of the Consumer Privacy Project at Stanford Law School*
11. **Brian Thomas,** *Managing Director at the Public Financial Management Group*

Appendix:
Initial Hypotheses – Analysis

The biggest obstructions toward autonomous driving technology being integrated into commercial personal vehicles are not technological but rather result from government regulation and liability issues.

It turns out that this hypothesis was not exactly correct. Right now, the level of the technology simply has not caught up to the expectations of the commercial market. The "mean failure distance" for autonomous vehicles is only on the order of 1000 miles. At this level of reliability, drivers may be caught off guard in the somewhat rare but expected event that they need to intervene. Currently, it is also impossible to design a good warning system for when intervention is necessary.

This is not to say that governmental and liability related issues are nonexistent. These concerns are present as well, but cannot truly be addressed until the technology has reached a point that makes it commercially viable.

The introduction of autonomously driven vehicles would need to occur in stages wherein the luxury car market would be the first target. This process would be aided by the help of government promotions—such as express lanes dedicated to autonomous vehicle use.

While the luxury car market has started to adopt some of the component technologies associated with autonomous driving, this is not the same thing as offering autonomous vehicles. The stages of implementation for the technology may have more to do with components than with tiers of automakers based on price. By this we mean that each "stage" would consist of a new piece of full autonomy being implemented in different vehicles. Since these pieces are mostly about safety and less about convenience or luxury, they could be implemented anywhere in the passenger vehicle marketplace as long as there are customers willing to pay for them.

It turns out that support through government incentives is highly unlikely. However, public-private partnership agreements could potentially be worked out to create necessary infrastructure; this is so long as these activities do not stir up too much political trouble.

The mainstream implementation of automatic driving technologies would potentially change the culture of driving on a worldwide scale.

This is true without a doubt, but there may exist considerable backlash to uniform implementation of autonomous vehicles. For the foreseeable future at least, autonomous vehicles will have to find a way to share the road with us human drivers—no matter how unpredictable we are.