

A reprint from

American Scientist

the magazine of Sigma Xi, The Scientific Research Society

This reprint is provided for personal and noncommercial use. For any other use, please send a request Brian Hayes by electronic mail to bhayes@amsci.org.

Leave the Driving to It

Brian Hayes

JANE HAS A MEETING this morning, so the car comes to pick her up at 8:15. En route, she finishes her breakfast, reviews her PowerPoint slides, updates her Facebook status, and does her daily KenKen. After the car delivers her to the office, it drives to a parking garage on the outskirts of the city, where it slips into a low, narrow slot. Later it will take young Judy and Elroy to their music lessons, then stop for a load of groceries before bringing Jane home. The car also has an errand of its own on today's agenda: the quarterly inspection and recertification required of all licensed autonomous vehicles.

Cars that drive themselves were already a cliché of futurist fantasies 50 years ago, and their long association with cartoonish fiction and dioramas at the World's Fair makes it hard to take the idea seriously. Nevertheless, sober thinkers believe it may be only a decade or two before the family car has a computer in the driver's seat. It's not too soon to ponder the social, economic and cultural consequences of such a development.

Already, some cars come equipped with "driver assistive technologies." There's adaptive cruise control, which keeps an eye on the car ahead and maintains a steady separation. Another system warns the driver if the car begins to stray outside its proper lane. And a few models even offer hands-free parallel parking.

More ambitious levels of automation are at the research-and-testing stage. In 1997 eight cars paraded down a San Diego freeway with the drivers waving both hands out the window, like kids showing off on a roller coaster. That demonstration was per-

*How would lives
and landscapes
change if every car
had a computer
in the driver's seat?*

formed on a lane studded with magnetic markers to guide the vehicles, but more recent trials have not required such aids. In contests sponsored by the U.S. military, autonomous vehicles have successfully traversed rough terrain and dodged city traffic. Google has built a fleet of seven computer-controlled cars, which have driven 140,000 miles on public roads (with a human driver present but seldom intervening). And last summer four mostly driverless electric-powered minivans completed a 15,000-kilometer trip from Italy to China.

Even with these milestones behind us, huge challenges remain. Later in this column I'll return to those scientific and engineering obstacles, but first I want to play the what-if game. If we could put a cybercar in every garage, how would that change the rhythms and routines of daily life?

Car Culture

The automobile itself has already worked quite a radical transformation on the human environment. In the 19th century, cities were full of horses. New York had more than 100,000 horses stabled on the island of Manhattan, and vast meadows in the outer boroughs were mown to supply the animals with hay. Caring for horses, hauling their food and manure, driving horse-drawn wagons and carriages—these were oc-

cupations that formed a major segment of the urban economy. All this equine infrastructure was swept away by the advent of mechanized transport—first the electric street car, then the automobile, finally the truck.

Of course cars and trucks gave rise to a new and larger infrastructure of their own, encompassing everything from parking meters and traffic signals to the entire worldwide petroleum industry. New York lost its 100,000 horse stalls and got millions of automobile parking places. Cities everywhere have grown fairy rings of suburbs and exurbs. A lacework of highways knits together distant communities, while sometimes dividing nearby ones. In most parts of North America, the automobile is how people get to work, how they get away on vacation and how they get home for the holidays. And a car is not just a means of transport; for many of us it's also a medium of self-expression—both *what* you drive and *how* you drive make a public statement. For adolescents, getting a driver's license is an important rite of passage and a step toward emancipation.

A hundred years ago, when sales of the Ford Model T were approaching 100,000 a year, an astute observer might have been able to foresee some of these developments—if not the specifics of Levittown, McDonalds, Walmart, Holiday Inn, NASCAR and Jiffy Lube, then at least the general trend toward a culture shaped by and dependent upon automobiles. If society eventually shifts to computer-driven vehicles, many elements of that culture will have to readjust. The details are inscrutable, just as they were a century ago. What effect will driverless cars have on pizza-delivery services or on the speed-trap revenues of small towns in the Midwest? As yet there are no quantitative answers to such questions. Still, the future is not completely opaque.

Brian Hayes is senior writer for American Scientist. Additional material related to the Computing Science column appears at <http://bit-player.org>. Address: 11 Chandler St. #2, Somerville, MA 02144. E-mail: brian@bit-player.org



Two driverless automobiles, navigating under computer control, arrive at an intersection in the DARPA Urban Challenge, a contest staged in 2007 by the U.S. Defense Advanced Research Projects Agency. In the foreground is Junior, a vehicle built by a team at Stanford University (and emblazoned with the insignia of its other sponsors). The other vehicle is Odin, from team VictorTango at Virginia Tech, which has the right of way and is proceeding through the intersection. Both vehicles have roof racks bearing multiple cameras, radars and laser scanners, as well as other sensory systems mounted elsewhere. Among 11 finalists, six finished the assigned course. The winner was a team from Carnegie Mellon University; Stanford came in second and Virginia Tech third. The event was held on the streets of George Air Force Base, a decommissioned facility in southern California. (Image courtesy of the Defense Advanced Research Projects Agency.)

An Abundance of Caution

There's one prediction about driverless cars that I can make with confidence: If millions of them ever roam the public highways, they will be far safer than cars driven by people. My confidence in this assertion does not derive from mere faith in technology. It's just that if robotic drivers were as dangerous as human ones, then computer-controlled cars would never be allowed on the roads. We hold our machines to a higher standard than ourselves.

Over the past decade, the number of auto accidents in the United States—counting only those serious enough to be reported to the police—has been running at about six million a year. Those accidents kill about 40,000 people and injure well over two million more. Estimates of the economic impact are in

the neighborhood of \$200 billion. Much of that cost is shared among car owners through premiums for auto insurance.

This safety record certainly leaves ample room for improvement. An appropriate goal for automated vehicles might be to reduce highway carnage to the same order of magnitude experienced in other modes of transport, such as railroads and commercial aviation. That would mean bringing road fatalities down to roughly 1 percent of their current level—from 40,000 deaths per year to 400. (In terms of deaths per passenger mile, cars would then be the safest of all vehicles.)

Could automation achieve such a hundredfold improvement? Taking the controls out of human hands would eliminate several major causes of crashes: drivers who are asleep, inebriated,

impatient, inattentive, overconfident, inexperienced. Unfortunately, computer drivers have foibles of their own; at the moment, they cannot even equal human performance, much less surpass it. I believe these failings can be overcome (and I'll explain how below). But for now let's just make the assumption that car travel can indeed be made safe, and examine the consequences.

Among young people in the United States, auto accidents are the leading cause of death; between ages 15 and 24, a third of all deaths happen on the road. For this age group, then, safer cars would have a noticeable demographic impact. The financial impact would also be significant—saving \$200 billion a year in accident losses, shutting down a \$200 billion industry that repairs smashed cars and bodies.

If accidents become rare enough, one would expect to see changes in attitudes and behavior, and perhaps in the design of vehicles. Current custom insists that we always buckle our shoulder harness and strap down the children in the back seat, as if every trip to the grocery store might end with a wreck—as indeed it might. No such precautions are taken on trains, subways, trolleys and buses. As accident rates fall, perhaps we would relax our vigilance in the car as well. And if the driver gradually becomes just another passenger, the interior of the automobile could look less like a padded cell or an airplane cockpit and more like an office, a theater, a café or even a bedroom.

A 99-percent improvement in safety will still leave tens of thousands of auto accidents every year, which may require new legal and financial mechanisms for compensating victims. Most collisions today are attributed to driver error rather than a defect or malfunction in the vehicle. When the vehicle *is* the driver, this distinction is no longer meaningful. Car manufacturers might be held liable for a larger share of the accidents—a responsibility they are certain to resist. (A legal

analysis by Nidhi Kalra and her colleagues at the RAND Corporation suggests this problem is not insuperable.)

The kinds of accidents caused by car-driving computers might be quite different from common human goofs. Of special concern is the possibility of a design flaw in hardware or software that could affect many cars the same way. It's easy to imagine a scenario in which one car after another follows exactly the same trajectory off a cliff.

Another essential point is that safety has to be balanced against reliability and robustness. Accidents can be reduced to near zero by adopting sufficiently conservative rules of operation. But nobody wants a car that pulls off the road and shuts down when the first snowflake falls.

The Road More Traveled

Almost 90 percent of American workers commute by car, most of them alone, with a median trip duration of just under half an hour each way. Although some people report that they enjoy this daily respite between work and home, many others find it tedious, and on the whole it seems a waste of human potential to spend so much time merely supervising the operation of a machine.

If we didn't have to keep our hands on the wheel and our eyes on the road, the time might be put to better use.

Other resources besides human attention might also be used more efficiently in a world of automated driving. Roadway real estate is one of them. Steven A. Shladover of the University of California, Berkeley, points out that only about 5 percent of the roadway surface is occupied by cars on a freeway running at peak throughput conditions (about 2,200 vehicles per lane per hour). Computer control, he suggests, could double or triple the density while keeping speeds constant. By squeezing more cars onto the same roads, we relieve pressure to widen highways or build new ones. Cars could be packed tighter both by narrowing the lanes and by reducing the headway between successive vehicles traveling in the same lane. Closer spacing is made possible in part because computers can maintain more precise control, both laterally and longitudinally. Equally important, automated vehicles can coordinate their motions through car-to-car data links, communicating not only their present position and velocity but also their intentions, such as changing lanes.

Communication is the key to many of the most attractive features of automated driving. (Indeed, Shladover suggests that the term "autonomous vehicle" is misleading in this respect; cooperation is more important than autonomy.) Consider the stop sign and the traffic signal: By forcing drivers to take turns at an intersection, these devices ensure that everyone gets access to a shared resource. But the stop-and-start regime of city traffic also causes congestion, wastes fuel and frays nerves. If cars could communicate with one another and with the roadway infrastructure, they could negotiate priority as they approached each intersection, adjusting their speeds to avoid conflict. Ideally, no car would ever have to make a full stop; traffic management would become a kind of precision choreography.

Paving Paradise

One aspect of car culture that would have been hard to foresee in 1911 is the extraordinary importance of parking. Although estimates vary, it seems there are two or three parking spaces for every car in the United States; think of it as one space at home, one at work, and a share of one at the mall. At just



A "Look ma, no hands!" demonstration of automated driving took place in 1997 in San Diego, where eight Buick sedans cruised down Interstate 15 in close formation, with some of the "drivers" waving out the window to show they were not in control. The test cars are on a closed roadway (ordinarily used for high-occupancy-vehicle traffic); to help the cars stay in their lane, sensors on the underbody detect magnetic studs installed in the concrete. The demonstration was part of a program, sponsored by the U.S. Department of Transportation, that was canceled a year later. (Image courtesy of the National Automated Highway System Consortium.)

two spaces per car, this works out to 500 million spaces in all, covering a total area of at least 3,000 square miles. A suburban shopping mall dedicates more land area to parking than to retail space. In dense central cities, parking spaces become a crucial limiting resource; they are the throttle valve that determines how many cars can come downtown. A study by Donald Shoup and his colleagues at the University of California, Los Angeles, shows that 30 percent of the drivers in some business districts are driving in circles, searching for a parking place.

The solution suggested by the cyber-car fantasy is automated valet parking: The car drops you off at the front door, then drives away to find a parking place on its own. When you're ready to leave, you call for the car and it comes to fetch you. In this scheme, the parking facility need not be within walking distance of your destination, so it can be moved out of the central business district. Likewise the suburban mall no longer has to be an island in a vast sea of asphalt. Another advantage of having cars park themselves is that they can squeeze together more tightly. (For one thing, there's no need to leave room for opening doors.)

All these strategies for making car trips more convenient and less tedious have to be seen as an inducement to ever-greater reliance on private vehicles. People will tolerate longer commutes if they can read or nap while in transit. They'll take the car to the city if they don't have to worry about where to park it. Thus a likely result of automotive automation is further diffusion of population over the landscape.

Another seemingly inevitable effect of more private vehicles is less public transit. However, all the advantages of computer-driven vehicles are also available to transit operators. Indeed, it's even possible that the distinction between public and private vehicles would blur a little. In a world where cars drive themselves and come to you when beckoned, there's not a lot of difference between calling for your own car and calling for a shared vehicle. For that matter, in a world where cars drive themselves and know their way around, there's not a lot of difference between a taxi and a rental car. Short-term rental programs such as Zipcar might thrive in this environment.

The effects that driverless technology would have on energy consump-

tion and carbon-dioxide emissions are hard to gauge. Given the same basic engine and vehicle, replacing a human driver with a computer ought to improve fuel economy. Platoons of closely spaced cars traveling together also save fuel through aerodynamic efficiency, although the effect may be small except at NASCAR speed. More important is the ability of cooperative scheduling and traffic management to avoid needless braking and acceleration. All of these factors are encouraging, but if the technology leads to more trips and longer trips, gains could turn to losses. And then there's the issue of all those empty ghost cars shuttling back and forth to remote parking lots or running errands on behalf of their owners. The average occupancy of vehicles on the road could fall below one person.

Keep on Truckin'

One group who will not be keen on driverless vehicles are those who drive for a living. Roughly four million people in the United States work as drivers of cars, buses and trucks—about 2½ percent of the workforce. If autonomous vehicles become commonplace, all of those jobs are in jeopardy. That's a major economic disruption.

Cost incentives might well lead long-haul trucking companies to become early adopters of driverless technology. Even if computer control were to double or triple the initial cost of the vehicle, this capital expenditure would still be far less than total payments to drivers. And a driverless truck does not have to stop for meals or sleep, or for anything other than fuel. Whereas a coast-to-coast trip with a solo human driver takes four days (obeying speed limits and regulations on hours of work), an automated truck could cover the same route in two days.

Taxi service presents similar issues. Eliminating the driver would dramatically lower costs for both the operator and the customer. (Presumably, the computer would not even expect a tip.)

Apart from professional drivers, other constituencies might also resist and resent autonomous vehicles. Car enthusiasts are an example: What's the point of owning a Ferrari if only the computer gets to drive it? Pedestrians and bicyclists would also be displaced if some roads were open only to automated vehicles.

Meanwhile, driverless technology would surely be welcomed by another

large group: people who cannot drive because of age or disability. An interesting subcategory is children, who are often dependent on their parents for transportation. Suppose we had a safe and reliable car that could be programmed to drive to any destination. Would we feel comfortable loading the kids into the back seat and sending them off to grandpop's house unaccompanied? I put this question to the parents of my own grandchildren (ages 5 to 12). The answer was no—and the children concurred. So it appears the self-driving car will not bring the liberation of the soccer mom.

The Open Road

The automobile has long been an emblem of personal liberty, particularly in American culture. With a car you can go where you want, whenever you want. Paradoxically, though, owning and operating a car is the most heavily regulated aspect of modern life. The car has to be registered and inspected; the driver has to be licensed; both of them have to be insured. You're welcome to put the top down and go for a spin, but you had better obey the speed limit and the stoplights.

The layers of regulation would surely get thicker with computer-driven cars. When platoons of automated vehicles are driving in tight formation, weaving through cross-traffic on a millisecond schedule, every car has to trust all the rest to behave predictably. Before being allowed to join the traffic stream, each car would have to provide some assurance that its hardware and software are functioning correctly and have not been tampered with; this might be done through a cryptographic authentication protocol. A side effect is that tinkering with your car, except for the most superficial changes, would likely be forbidden.

Road transport might become more like the airline system. Major airports enforce a scheme in which flights are not allowed to depart until they have secured a landing "slot" at their destination—a place in the sequence of anticipated arrivals. An analogous rule could regulate access to congested highways or bridges; you couldn't leave home until the road had room for you. The reward for submitting to this regimen of regulation would be faster and more predictable travel.

The operation of such a traffic-control system raises fascinating questions of

social choice. Consider the case of cities connected by two parallel roads. Travelers free to choose their own route will presumably always take the road that minimizes their own travel time. But this exercise of free choice can contribute to congestion that delays everyone. An algorithm may be able to assign cars to roads in a way that outperforms the sum of everyone's selfish choices. Would travelers accept such interference with their liberties? And how do we agree on the optimal outcome? What if a plan speeds up 99 percent of the cars, but the remaining 1 percent suffer a two-hour delay?

Are We Almost There Yet?

The car that will take Jane to the city and then go park itself will not be a 2012 model. The car that can reduce highway fatalities by 99 percent will not be a 2012 model either. Measured against those goals, the current state of the art for robotic vehicles looks pretty wimpy—but not hopeless.

Here's a story that gives a hint of what computer-controlled driving is like today. A vehicle named Talos, built by a team at MIT, was being prepared for the DARPA Urban Challenge, a competition staged in 2007 by the Defense Advanced Research Projects Agency. In preliminary testing, Talos was nearing an intersection when it noticed another vehicle approaching on the cross street. The other car had the right of way, and so Talos stopped to wait. But then, because of "sensor noise," Talos momentarily lost track of the car. When the sensors reacquired the signal, the planning and guidance logic concluded it was seeing a new vehicle, which had just arrived at the intersection. This meant that Talos had priority, and so it started forward. So did the other car. They didn't collide, but it was a near thing.

This incident offers a fascinating glimpse into the mind of an alien intelligence, trying to make sense of a world where cars can pop in and out of existence without warning. The story also suggests we are still a long way from creating a computational agent with the kind of common sense needed to pass the road test for a driver's license. I believe this assessment of the situation is correct, but it gives the wrong impression about the prospects for building driverless automobiles.

Much of the recent work on autonomous vehicles treats driving as a

problem in artificial intelligence and computer vision. The challenge is to extract meaning from sensory signals and form an accurate conceptual model of the roadway situation. With great effort, this approach may eventually succeed, though perhaps only in creating a computer driver as fallible as a human one.

There's another way to go about it. Instead of trying to replicate the driver's sensory faculties and mental model of the world, we can reengineer the world itself so that sophisticated perception and cognition are no longer needed. Consider again two cars at an intersection, trying to decide who goes first. Human drivers rely on subtle forms of communication, including eye contact and occasional hand-waving, to resolve this situation. Perhaps a computer could learn to do the same thing, but an easier course is to provide a data channel over which the cars can communicate and negotiate directly.

Aviation offers a useful point of reference. Commercial aircraft routinely fly under computer control (except at takeoff and landing). But the aircraft autopilot does not look out the window and try to interpret visual cues. Instead, the flight-management system relies on ground-based beacons, satellite signals and inertial navigation, as well as plane-to-plane data links for collision avoidance. Meanwhile a central facility (air-traffic control) coordinates the movements of aircraft and resolves conflicts.

Admittedly, navigation and traffic management are easier in the wide-open spaces of the sky than on crowded, quasi-one-dimensional roadways, but the same principles could be applied. With appropriate infrastructure, each car could have accurate and timely information about the state of the roadway and nearby cars—their positions, velocities and intentions. Even with such complete information, computing optimal paths for all the cars remains formidably difficult, but it is a problem of algorithms and control theory, not cognitive science.

The key first step in making this approach feasible is building a communication network linking nearby vehicles and roadside relay stations. Standards for such networks are already in preparation. Some version of the network is likely to be implemented soon for less-grandiose purposes, such as traffic reporting and entertainment.

Building new infrastructure is slow and expensive, so even if all the technological problems were solved, it would be years before large numbers of cars on large numbers of roads were routinely taking charge of their own movements. Thus there's time for planning and choosing what kind of transportation system we'd like to have. It's a good moment to be asking not just "Can we get there from here?" but also "Do we want to go?"

Bibliography

- Broggi, A. 1999. *Automatic Vehicle Guidance: The Experience of the ARGO Autonomous Vehicle*. River Edge, N.J.: World Scientific.
- Campbell, M., et al. 2010. Autonomous driving in urban environments: approaches, lessons and challenges. *Philosophical Transactions of the Royal Society* 368:4649–4672.
- Chester, M., A. Horvath and S. Madanat. 2010. Parking infrastructure: energy, emissions, and automobile life-cycle environmental accounting. *Environmental Research Letters* 5:034001.
- Fletcher, L., et al. 2008. The MIT–Cornell collision and why it happened. *Journal of Field Robotics* 25:775–807.
- Folsom, T. C. 2011. Social ramifications of autonomous urban land vehicles. IEEE International Symposium on Technology and Society, May 2011, Chicago. http://www.enviroteach.com/social_ramifications.pdf.
- Furda, A., and L. Vlacic. 2011. Enabling safe autonomous driving in real-world city traffic using multiple criteria decision making. *IEEE Intelligent Transportation Systems Magazine* 3(1):4–17.
- Hartenstein, H., and K. P. Laberteaux. 2010. *VANET: Vehicular Applications and Inter-Networking Technologies*. Chichester, U.K.: John Wiley and Sons.
- Kalra, N., J. Anderson and M. Wachs. 2009. Liability and regulation of autonomous vehicle technologies. California Partners for Advanced Transit and Highways Research Report UCB-ITS-PRR-2009-28.
- McCarthy, N. 2009. *Autonomous Systems: Social, Legal and Ethical Issues*. London: Royal Academy of Engineering. <http://www.raeng.org.uk/autonomoussystemts>.
- Ni, D., et al. 2010. Preliminary estimate of highway capacity benefit attainable with IntelliDrive technologies. In *Proceedings of the 13th International IEEE Conference on Intelligent Transportation Systems*, pp. 819–824.
- Schakel, W. J., B. van Arem and B. D. Netten. 2010. Effects of cooperative adaptive cruise control on traffic flow stability. In *Proceedings of the 13th International IEEE Conference on Intelligent Transportation Systems*, pp. 759–764.
- Shladover, S. E. 2009. Cooperative (rather than autonomous) vehicle-highway automation systems. *IEEE Intelligent Transportation Systems Magazine* 1(1):10–19.
- Shoup, D. 2007. Cruising for parking. *Access: The Magazine of the University of California Transportation Center* 30:16–22.
- Thrun, S. 2010. Toward robotic cars. *Communications of the ACM* 53(4):99–106.