Editor: Daniel Zeng
University of Arizona and
Chinese Academy of Science
zeng@email.arizona.edu

Self-Driving Cars and the Urban Challenge

Chris Urmson and William "Red" Whittaker, *Carnegie Mellon University*

Self-driving cars have been a dream as long automobiles have existed. The automobile is ubiquitous in the developed world and is becoming so in the developing world. In 2007, the world's two largest automakers sold over 18 million vehicles worldwide. As we consider domains to which we can apply intelligent systems, the automotive industry stands out as having the most potential for impact.

Benefits of autonomous vehicles

Each year more than 1.2 million people die in traffic accidents. In the United States, over 42,000 people are killed and over 2.7 million people are injured each year.[1] The prevailing belief in the automotive industry is that the benefit of passive safety systems such as seat belts and air bags has reached a plateau. To improve safety, vehicles must avoid crashes rather than attempt to survive them.

Autonomous vehicles provide more benefits than just safety. Each year, the average American adult spends 100 hours commuting, time that could otherwise be spent relaxing, reading, or even working. For individuals who have lost their driving privilege owing to age or disability, autonomous vehicles will remove the spatial barriers limiting them from interacting with family, friends, and the community. This not only makes sense for society but also commercially since it increases the number of people who would consider purchasing automobiles.

Autonomous vehicles will also reduce the environmental impact of driving. Autonomous cars and trucks can tune their acceleration and deceleration profiles to reduce wasted fuel. As autonomy technology becomes ubiquitous, manufacturers will be able to reduce vehicle mass because much of that mass is devoted to protecting the occupants. Because there's a linear relationship between vehicle mass and fuel consumption,[2] it will also be possible to trade intelligence for vehicle mass, increasing both safety and fuel efficiency.

A brief history

These benefits have led to numerous research programs around the world. Much early research in autonomous vehicles focused on highway driving. Through the '80s and '90s, Ernst Dickmanns, Charles Thorpe, and Alberto Broggi led the most visible programs in highway driving.[3] In the '80s, Dickmanns' van was able to drive on roads without other traffic. By the late '90s, the Navlab, VamP, and ARGO teams demonstrated thousands of kilometers of highway driving at speeds up to 130 kph with normal traffic. For these teams, the maximum distance for a single autonomous drive was on the order of 100 km.

Since these projects, interest has increased in vehicles that can operate in the more complicated problem domain of urban driving. One example is the e-Motion group, which has developed the CyCab,[4] capable of driving autonomously at a low speed among pedestrians. This program aims to provide personalized mass transit for dense urban environments. While these vehicles are becoming a viable alternative to other forms of public transit, they still can only travel at low speeds.

In 2003, DARPA announced the first Grand Challenge. Although no vehicle was able to complete this challenge, the competitors set a new benchmark for autonomous capability and provided a template on how to complete it. The next year, five vehicles completed a similar challenge, with Stanley[5] edging out Sandstorm and H1ghlander[6] to complete the 152-mile race in a little under seven hours.

As a next step, DARPA organized the Urban Challenge. The challenge called for autonomous vehicles to drive 60 miles through an urban environment, interacting with other moving vehicles and obeying the *California Driver Handbook*.[7]

Boss and the Urban Challenge

Boss, a modified 2007 Chevy Tahoe (see figure 1), won this challenge using a combination of laser, radar, and GPS data to safely navigate a 53-mile test course among 60 other vehicles (10 autonomous and 50 human-driven).

1541-1672/08/$25.00 © 2008 IEEE
Published by the IEEE Computer Society

The challenge was announced in May 2006 and took place in November 2007. So, teams had only 18 months to research, design, and implement a solution. It quickly became apparent that a key to completing the challenge was to determine the bounds of what was called for and focus only on solving those problems.

The rules for the challenge provided some guidance.[8] Only vehicles would be on the course, and they had to be midsized cars or larger. This removed pedestrians and bicycles from the scope. The only traffic controls on the course would be stop signs, and the location of the corresponding stop-lines would be provided. This eliminated traffic lights, yield signs, and the need to detect or read them from the scope. The rules also indicated that the roads a vehicle could drive on would be at least partially defined by highly accurate GPS waypoints. This became a key rule that participants could utilize to improve system performance by reducing complexity.

Accurate but sparse GPS data can be combined with aerial imagery to provide accurate, dense definitions of road geometry. Once a framework of dense road geometry is available, the problem of road navigation can be transformed from estimating a road's shape to estimating a position relative to a road: a much lower-dimensional and thus easier problem. Boss uses this approach. Instead of emphasizing a forward-looking vision system to extract the road shape, it uses a pair of down-looking lasers mounted above the rear wheels to detect lane markings. A filter combines the measured location of lane markings, a map, and GPS/INS (inertial navigation system) data to estimate Boss's position. Once Boss knows its position on the road, it can predict upcoming corners, slowing and turning as necessary on the basis of map information.

Maps also provide important contextual information. Consider figure 2a; without any context, the two vehicles appear to be on a collision course. Given context (see figure 2b), we can safely assume that one of the vehicles is just making a lane change and that nothing untoward is about to happen. Boss uses a model of the nearby roads in much the same way to anticipate and reason about other vehicles' movements.

Beyond basic driving, Boss can park, pass stopped and moving vehicles, and safely interact with other vehicles at intersections while driving at up to 30 mph.

Figure 1. Boss, a modified Chevy Tahoe. Boss completed the 2007 DARPA Urban Challenge, safely navigating 53 miles through an urban environment among other moving vehicles.

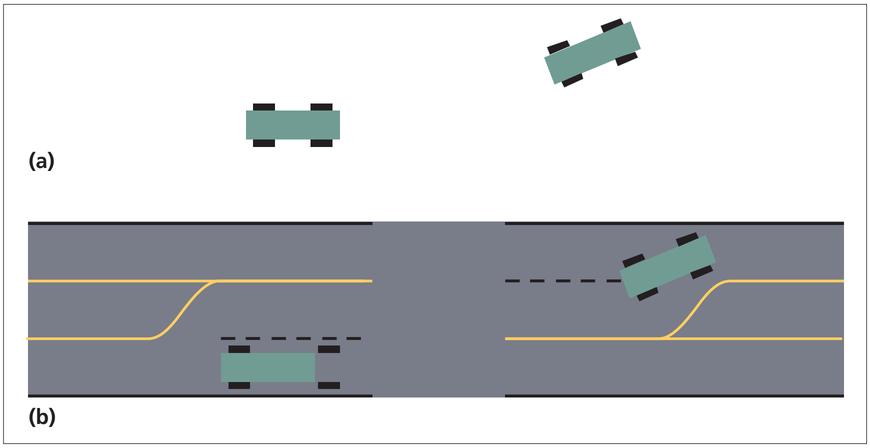

Figure 2. The importance of context for urban driving: (a) Without any context, the two vehicles appear to be on a collision course. (b) Given the context, we can safely assume that one of the vehicles is just making a lane change.

Implementing these capabilities requires a complex software system. Our system has five main components: perception, mission planning, motion planning, infrastructure, and behavioral reasoning. The perception subsystem models other moving vehicles and static obstacles around Boss, along with localizing Boss relative to the road model. The mission-planning subsystem develops the best route between two locations on the map, and reroutes Boss when a road is blocked. The motion-planning subsystem drives Boss and ensures that it doesn't collide with anything. The software infrastructure provides the necessary glue that lets each subsystem communicate and enables efficient testing and development through data logging, playback, and configuration management. Finally, the behavioral-reasoning subsystem makes tactical decisions, executing the mission plan by dispatching goals to the motion-planning subsystem, reasoning about intersections and other vehicles, and handling recovery from error states.

One of the key tenets of Boss's software system is to never give up. Boss's error recovery system enables it to always attempt some maneuver. The recovery system acts in a context-specific manner, with different error recovery strategies based on whether Boss is currently driving on a road or maneuvering through a parking lot, and so on. All the recovery modes share some aspects: they attempt increasingly risky maneuvers as time progresses, and they generate a nonrepeating series of motion goals. Despite significant testing, the error recovery system was invoked several times during the Urban Challenge final event and was thus critical to Boss's success.

During the testing, qualifications, and final event, Boss proved itself to be a capable vehicle. The team's intense test schedule included over 3,000 km of autonomous operations in test sites in three different states. During qualifications, Boss consistently performed well, completing each evaluation course. The final event provided the most challenge; Boss had to resort to using its error recovery system 11 times, but in the end finished about 20 minutes faster than its closest competitor. Further details of how the system worked and how the team tested the vehicle appear elsewhere.[9]

Beyond the Urban Challenge

Although the Urban Challenge was a resounding success, the challenge of fully autonomous urban driving hasn't yet been met. Despite the Urban Challenge vehicles' strengths, none of them could interact with traffic lights, most would operate poorly around pedestrians, and all relied on sensors that are too expensive or unwieldy for consumer vehicles. As the bounds of what's possible for autonomous vehicles are pushed forward, these technical issues and the social and legal issues associated with relinquishing driving control will be at the forefront of the research agenda. Despite the challenges, the future appears to be now. Heavy-equipment manufacturers are announcing autonomous haul trucks, with some ready for deployment as early as 2010. General Motors announced that autonomous vehicles will be ready for market by 2020. These are heady times for autonomous-vehicle research. ■

References

1. *Traffic Safety Fact Sheet*, US Nat'l Highway Traffic Safety Administration, 2005.
2. *Final Economic Assessment, Corporate Average Fuel Economy Standards for MY 2005–2007 Light Trucks, Section V*, US Nat'l Highway Traffic Safety Administration, May 2003.
3. M. Bertozzi et al., "Vision-Based Intelligent Vehicles: State of the Art and Perspectives," *J. Robotics and Autonomous Systems*, vol. 32, no. 1, 2000, pp. 1–16.
4. C. Pradalier et al., "The CyCab: A Car-Like Robot Navigating Autonomously and Safely among Pedestrians," *Robotics and Autonomous Systems*, vol. 50, no. 1, 2005, pp. 51–68.
5. S. Thrun et al., "Stanley, the Robot That Won the DARPA Grand Challenge," *J. Field Robotics*, vol. 23, no. 9, 2006, pp. 661–692.
6. C. Urmson et al., "A Robust Approach to High-Speed Navigation for Unrehearsed Desert Terrain," *J. Field Robotics*, vol. 23, no. 8, 2006, pp. 467–508.
7. *California Driver Handbook 2008*, Calif. Dept. of Motor Vehicles, 2008, www.dmv.ca.gov/pubs/dl600.pdf.
8. *Urban Challenge Rules*, DARPA, 2007, www.darpa.mil/grandchallenge/docs/Urban_Challenge_Rules_102707.pdf.
9. C. Urmson et al., "Autonomous Driving in Urban Environments: Boss and the Urban Challenge," submitted to *J. Field Robotics*, 2008.

Chris Urmson is the Director of Technology for the Urban Challenge at Carnegie Mellon University's Robotics Institute. Contact him at curmson@ri.cmu.edu.

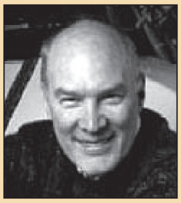

William "Red" Whittaker is the Fredkin Professor of Robotics at Carnegie Mellon University's Robotics Institute. Contact him at red@ri.cmu.edu.