Progress in Self-Driving Vehicles

Chris Urmson, Google

Automated driving has experienced a research renaissance over the last decade. As a research community, we have been motivated by the opportunity to increase safety, increase mobility, and improve the experience of mobility. Some of the key advancements that have shaped the field over this time period have been the advancement and application of machine learning, advancements in large scale mapping, improved LIDAR and RADAR sensing capability, and more recently, a deeper understanding of the human factors that will influence the form by which this technology comes to market.

Why Self-Driving Vehicles?

In the United State, the leading cause of death for individuals aged 4-34 is traffic accidents (Hoyert, 2012). We kill over 30,000 people each year on our roads, and 90+% of these accidents are due to human error. The importance of personal mobility in our society is such that when individuals lose the privilege of driving, and lose social connections it enables, their life expectancy drops precipitously (Edwards, 2009). The ability to move through cities is decreasing as more and more users, longing for individual automobile mobility, flood roadways. The rate of urbanization in developing cities is the latest incarnation of the tragedy of the commons.

Self-Driving vehicles offer the promise of addressing all of these challenges – they should dramatically reduce accidents, enable people who cannot drive to get around, and when deployed as part of an efficient shared vehicle fleet, reduce congestion. The idea is not new, and the current highly visible efforts build on a deep foundation of technical excellence.

A Deep History

The history of self-driving vehicles is long. As early as the 1939 World's Fair, GM showed a concept of the automated roadway of the future. In 1950, GM R&D introduced the Firebird II concept car, capable of following buried cables that emitted an RF signal. During the 80's and 90's the introduction of the micro-computer enabled practical, on-line computation on a mobile platform. Ernst Dickmanns was a pioneer in this space, introducing early versions of foveated stereo-vision systems (Dickmanns, 2007). In the mid 90's machine learning began to be applied to the problem. RALPH (Thorpe, 1990) (and related work) was one of the earliest applications of machine learning (neural networks in this case) applied to automated driving. The combination of RALPH (combined with a nascent forward looking RADAR system) enabled vehicles to drive thousands of miles in 1997. Elements of this technology have found their way into lane keep assist systems and forward collision mitigation braking and adaptive cruise control systems.

The Grand Challenges

Much of the on-road automated driving work faded after the successful 1997 National Automated Highway Systems Consortium demonstration – the technology worked reasonably well, but automated driving research funding turned towards the military while the automotive industry slowly commercialized driver assistance systems. In 2003, the driving research community was re-energized by the announcement of the DARPA Grand Challenges. The Floyd D. Spence National Defense Authorization Act for Fiscal Year 2001, called for 1/3rd of all US Military ground vehicles to be unmanned by 2015. In a 2002 report, the National Academies indicated that this goal would not be achievable, and the Department of Defense should pursue other strategies for achieving this goal (Rose, 2002). Thus DARPA's Challenges were born.

Figure 1. Stanley (left), Sandstorm (right) and H1ghlander, were the top three finishers in the DARPA Grand Challenge.

The initial Grand Challenges were off-road races across the desert, with the notional goal of having autonomous vehicles drive from Los Angeles to Las Vegas without remote assistance. In 2004 the challengers went only 7 miles of the 150-mile course (Urmson, 2004). In the following year, several vehicles completed the competition, with a team from Stanford winning (Thrun, 2006). Despite the relatively short timeframe, several notable technical innovations were incorporated into the vehicles. All of the competitors were given a coarse map of the route, but several of the successful teams augmented the map data with information available from other publicly available sources - this notion of fusing past data in conjunction with onboard sensing data was a novel concept at the time (Urmson, 2006). The approach was enabled by newly available access to high resolution aerial imagery, and gave the vehicles a degree of foreknowledge of the terrain that enabled, better and safer driving than had been demonstrated prior.

The Stanford team used machine learning techniques extensively. The vehicle used machine learning to bootstrap it's visual system using it's LIDAR sensors, to allow it to drive faster than was possible using LIDAR alone. It was able to detect when it encountered rough terrain, and slow appropriately using a learned model of "bumpiness". Their success in the challenge helped reinforce machine learning's value in the field of autonomous driving.

The Urban Challenge

While the Grand Challenge was indeed a grand challenge, the vehicles operated in a world devoid of other moving vehicles. Case in point – when Stanley, the Stanford vehicle, eventually passed H1ghlander, the Carnegie Mellon vehicle, to claim the victory, H1ghlander was paused, and Stanley passed an inert vehicle. The Urban Challenge was thus the next evolution of the competition, where the vehicles were now forced to not only complete the challenge with moving vehicles, but to obey a subset of the driving rules that human drivers take for granted (stay in the lane, obey precedence rules at intersections, avoid other vehicles, etc.). The competition was staged in 2007, with vehicles required to drive 60 miles around a decommissioned air force base in Victorville, California. At the end of the day, six vehicles finished the competition, with teams from Carnegie Mellon, Stanford and Virginia Tech rounding out the top three (Buehler, 2009).

The key technical advancements came in the form of high-density LIDAR and an increased demonstration of the value of high-density maps. The LIDAR sensors used in the Grand Challenge were single plane LIDARs , sometimes actuated to sweep volumes, but generally carefully calibrated to sweep scan lines through the environment as the vehicle moved. The Urban challenge introduced the concept of high-density LIDARs, through the sensor developed by Velodyne – a spinning sensor head that swept a set of 64 individual LIDAR emitters through space, generating over a million range measurements per second with relatively high angular resolution. This style of sensor enabled a new level of precision modelling that had been difficult, if not impossible, to achieve in real-time before.

The value of digital maps came to the forefront during the Urban Challenge. By utilizing the maps, vehicles were able to anticipate the likely trajectory of other vehicles and focus their attention in appropriate directions at intersections. Furthermore, by utilizing the map as a guide, vehicles were able to utilize their limited computation efficiently.

Post Challenges

In the seven years since the last Grand Challenge, industry has taken up the gauntlet of advancing self-driving technology. In 2009, Google started a program to develop self-driving vehicles. Over the last five years, Google's vehicles have driven over 700,000 miles autonomously on public roads. The technology being developed by Google builds upon many of the themes developed during the DARPA Challenges. The vehicles utilize high resolution maps (now being built at city scale) to help guide the onboard system's perception and planning behaviors. The vehicles utilize a combination of LIDAR , camera and RADAR sensors to provide a partially-redundant and multi-spectral model of the environment. The onboard software system leverages the hundreds of thousands of miles of driving data and machine learning techniques to predict the behavior of other road users.

In parallel with Google's efforts, the automotive industry is broadly engaged in the development of advanced driver assistance systems, with all of the major car companies and their suppliers developing varying degrees of automated driving. The largest difference between the approaches of the classical automotive companies and Google's approach is the degree to which the driver is engaged in the driving task. Google is currently developing vehicles that would be fully self-driving - only requiring a rider to tell the vehicle where to go. The automotive companies are primarily focused on delivering advanced driver assistance systems that require the driver to remain in the steering loop. The latter approach requires a smaller incremental technical step, but is challenged by various problems with driver attentiveness and skill atrophy (Llaneras, 2013).

Figure 2. Google's prototype fully self-driving vehilce.

In the coming years we will see advanced driver assistance systems and self-driving vehicles become common place, delivering on the long held promise of making our roadways safer and more convenient for all.

References

Hoyert, D. L., & Xu, J. (2012). Deaths: preliminary data for 2011. *National vital statistics reports*, *61*(6), 1-51.

Edwards, J. D., Perkins, M., Ross, L. A., & Reynolds, S. L. (2009). Driving status and three-year mortality among community-dwelling older adults. *The Journals of Gerontology Series A: Biological Sciences and Medical Sciences*, gln019.

Dickmanns, E. D., & Wünsche, H. J. (2007). *Dynamic vision for perception and control of motion* (pp. I-XCII). London: Springer.

Thorpe, C., & Kanade, T. (1990). *Vision and navigation*. Kluwer Academic Publishers.

Rose, M. F. (2002). Technology Development for Army Unmanned Ground Vehicles. *Committee on Army Unmanned Ground Vehicle Technology, Board on Army Science & Technology, National Research Council*.

Urmson, C., Anhalt, J., Clark, M., Galatali, T., Gonzalez, J. P., Gowdy, J., ... & Whittaker, W. L. (2004). High speed navigation of unrehearsed terrain: Red team technology for grand challenge 2004. *Robotics Institute, Carnegie Mellon University, Pittsburgh, PA, Tech. Rep. CMU-RI-04-37*.

Thrun, S., Montemerlo, M., Dahlkamp, H., Stavens, D., Aron, A., Diebel, J., ... & Mahoney, P. (2006). Stanley: The robot that won the DARPA Grand Challenge.*Journal of field Robotics*, *23*(9), 661-692.

Urmson, C., Ragusa, C., Ray, D., Anhalt, J., Bartz, D., Galatali, T., ... & Struble, J. (2006). A robust approach to high-speed navigation for unrehearsed desert terrain. *Journal of Field Robotics*, *23*(8), 467-508.

Buehler, M., Iagnemma, K., & Singh, S. (Eds.). (2009). *The DARPA urban challenge: autonomous vehicles in city traffic* (Vol. 56). springer.

Llaneras, R. E., Salinger, J., & Green, C. A. (2013). Human factors issues associated with limited ability autonomous driving systems: Drivers' allocation of visual attention to the forward roadway. In *Proceedings of the 7th International Driving Symposium on Human Factors in Driver Assessment, Training and Vehicle Design* (pp. 92-98).