

Autonomous Driving–5 Years after the Urban Challenge: The Anticipatory Vehicle as a Cyber-Physical System

Christian Berger¹ and Bernhard Rumpe²

¹Department of Computer Science and Engineering
Chalmers | University of Gothenburg
Sweden
christian.berger@chalmers.se

²Software Engineering
RWTH Aachen University
Germany
rumpe@se-rwth.de

Abstract: In November 2007 the international competition DARPA Urban Challenge took place on the former George Airforce Base in Victorville, California to significantly promote the research and development on autonomously driving vehicles for urban environments. In the final race only eleven out of initially 89 competitors participated and “Boss” from Carnegie Mellon University succeeded. This paper summarizes results of the research carried out by all finalists within the last five years after the competition and provides an outlook where further investigation especially for software engineering is now necessary to achieve the goal of driving safely and reliably through urban environments with an anticipatory vehicle for the mass-market.

1 Introduction

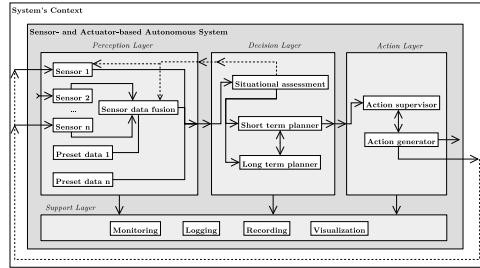
Autonomous driving was significantly fostered within that last decade due to three major challenges, which were carried out by DARPA—the research agency of the Department of Defense. The most exciting one was the 2007 DARPA Urban Challenge, in which competitors from all over the world had to develop robotic vehicles that were able to drive entirely without any human interaction in urban environments. Compared to the previous Grand Challenges from 2004 and 2005, autonomously driving vehicles had not only to deal with other moving vehicles but they also had to obey the Californian traffic law during all their autonomous operations.

In Fig. 1(a) the autonomously driving vehicle “Caroline” [RBL⁺08] developed from the Technische Universität Braunschweig under the organization of the authors is depicted. For that vehicle, the problem to drive without any human interaction was split into three main tasks according to Fig. 1(b): Data perception and preprocessing from all sensors to generate an environmental model, understanding this surroundings’ model to derive the next deriving decision, and performing actions within the system’s context i.e. steering and accelerating the vehicle. The results gathered during that competition meanwhile led to the development of her successor “Leonie” [NHO⁺11].

Current driverless vehicles have collected data from more than 140,000mi [Thr10a, Thr10b] or are already running in China [Nan11]. Although that competition was named “Urban



(a) Sensor setup of the autonomously driving vehicle “Caroline”.



(b) System architecture consisting of a data perception layer, a decision layer, and an action layer.

Figure 1: “Caroline”—the contribution from Technische Universität Braunschweig competing in the 2007 DARPA Urban Challenge and its principal system architecture.

Challenge” pedestrians and bicyclists had not been regarded at all. Thus, what has been achieved within the last five years after that spectacular final event and where are still aspects, which need more attention to further reduce traffic jams, to save fuel, and most important to prevent casualties and fatalities? This contribution summarizes research results from finalist teams from the last five years with a strong focus on aspects that were not explicitly included in the DARPA Urban Challenge competition.

2 Localization & Perceiving the Vehicle’s Surroundings

As shown in Fig. 1(a), a set of redundant sensors with overlapping viewing areas is necessary to create a reliable representation from the autonomously driving vehicle’s surroundings. Besides this sensors’ online data, preprocessed offline data improves the vehicle performance. [SU09] outlines an aerial image analysis, which relies on self-supervised machine learning to detect parking spots. From this analysis the fundamental topological structure for the area of interest is derived to provide an augmented graph as drivable paths in advance.

Due to limited accuracy or temporal lack of GPS, the precise position of the autonomously driving vehicle is fundamentally important. [LT10] as an extension to [LMT07] outlines an approach that models the static environment as a probabilistic grid with a $0.15 \times 0.15 \text{m}^2$ resolution where every cell represents a Gaussian distribution over infra-red remittance. Thus during online localization, an RMS-accuracy of nearly 0.1m could be achieved. However, the outlined implementation requires approximately 10MB/mi of storage and is not entirely independent from rough weather conditions or modifications to the environment (e.g. construction sites); moreover, elevation information could improve this map.

During the competition neither traffic lights nor traffic signs had to be detected. Nowadays, static traffic signs are provided as annotations within digital maps and the state of traffic lights can be provided wirelessly by vehicle-to-X. However for the latter, the infras-

structure itself must be modified, which would be very expensive (e.g. nearly 2,200 traffic lights exist in Berlin); an optical detection would cover already existing traffic lights. Thus, [LADT11] outlines a map-based approach to control the region of interest for detecting the position and orientation of a traffic light within camera images; [FU11] additionally encodes semantics (i.e. various flashing states) or labels like “dim” to indicate traffic lights that are difficult to detect in images to parametrize an algorithm. However, these approaches also unveil problems due to lens flares or heavy rain.

As depicted in Fig. 1(a), a static sensor set is mounted and thus, viewing angles are fixed and could not be adapted where necessary. In [SU08], a pointable sensor for covering unobserved areas of interest is outlined, which is adjusted by an information entropy model of an intersection’s uncertainty for example. This approach could be extended to model temporary occluded areas by e.g. dynamic objects. An alternative approach is outlined by [SP12] that proposes the combination of sensors with vehicle-to-X-communication.

The competition excluded explicitly pedestrians and bicyclists in general; however, for the broader usage of driverless vehicles these at least protected road users¹ must be detected reliably. Contrary to detecting pedestrians continuously within the driving area, [BCG⁺09] presents an approach that focuses on specific situations like stopped vehicles or crosswalks detected by a laser scanner. Within these pre-detected areas, pedestrians are validated by the computationally intense vision system. Contrary to pedestrians, bicyclists may also share the road with vehicles; in [CRZ11] a monocular vision-based approach is described, which uses a deformable shape model for detection and an interacting multiple model filter to track a bicyclist. In [WGR11b], a real-time detecting approach based on contour cues is outlined running at 20fps on VGA resolution and which could be improved by using a GPU. However, vision-based approaches depend on external illuminating and would be insufficient to operate reliably.

From a software engineering point of view, we will face the challenge that the mass-market will have sensors in many different forms and qualities. Furthermore, degradation of their quality as well as replacement by other sensors with different quality characteristics will be immanent. This poses special adaptivity to the software architecture that processes sensor data, both on easy adaptation and on self-awareness of the quality characteristics of the sensors in the car. Graceful degradation of software functions based on sudden failures of sensors need also to be handled reliably.

3 Understanding the Surroundings’ Data

The sensors perceive a vehicle’s surroundings, which must be analyzed and interpreted in real-time. In [HW11], a layer on top of the perception layer is introduced that abstracts details, which are currently irrelevant to the situation analysis and maps objects to a graph-based environmental model. Thus, deriving decisions is reduced to the objects, which are relevant to future actions like turning while yielding right of way to forthcoming vehicles. In [CWW11], an environmental graph-structure is also outlined that extended the road

¹According to Statistisches Bundesamt, 4,002 fatally injured road users (398 bicyclists) were counted in 2011.

graph used during the competition. Their representation includes annotations about way point interpolations, lane categories, and bridges or tunnels to be suitable for German autobahn and in urban environments [WGR11a].

As outlined in Sec. 2, aerial images could be preprocessed to provide maps for parking lots. In [SUWL10] these maps are used to generate proper drivable paths from an unstructured parking area. In contrast, [DTMD09] uses a Voronoi field to model path lengths and distances to obstacles for generating drivable trajectories. A more elaborated approach is described in [JHMBH10] that could also be applied when GPS is not available as in parking garages.

Reliable detection of pedestrians and bicyclists is crucial during the decision process. [ZRG⁺09] outlines an approach, which predicts their possible future trajectories. A similar focus is provided by [GSL11], in which plausible future motions for each tracked person are calculated. A more generic approach is outlined by [HC11] that uses hierarchical trajectory clustering to improve the computational effort. The goal is to cluster obstacle trajectories, which have an effect on the planned vehicle's trajectory.

[VGZD07] describes an approach for high-level situation reasoning that tries to predict the future evolution of a situation. The main idea is to assess the current traffic situation for finding a similar previous case to derive a proper decision. Hereby, a case is defined by the behavior of the autonomous vehicle, the behavior of other participants, and an estimation how similar the current traffic situation to that case is.

In [WDSL11], a Markov Decision Process is used to derive single-lane driving decisions; thus, uncertainties for a sensor's noise and the other vehicles' behavior can be modeled, which results in a robust driving behavior in uncertain situations compared to prior ACC systems. A more elaborated system for assisting the driver on highways is presented in [SFG⁺10]. This system evaluates continuously the vehicle's surroundings to warn the driver about unsafe lane change situations or to predict safe lane change speeds.

Recent research efforts also include data exchange realized by vehicle-to-X communication. [NHF⁺11] therefore proposes an approach that explicitly includes this data into a so called knowledge layer for deriving driving decisions. The main idea is to combine sensor-gathered information with wirelessly received information to generate a more reliable representation of the vehicle's surroundings. Thus, information from occluded situations non-visible to sensors can be considered.

From a software engineering point of view, the adaptivity on the sensor side will enforce high-adaptivity on the sensor data fusion and understanding. This enforces a reliable and flexible software architecture that processes data on several levels, but also understands to shortcut processing in case of emergencies (e.g. when suddenly braking is necessary). While the general architecture from Fig. 1(b) has been used in many projects, these architectures are generally not flexible enough and usually do not provide data connection shortcuts. Furthermore, many of today's experimental architectures for such an anticipatorily driving vehicle assume a rather centralized form of processing while today's car manufacturers' architectures use massively distributed control unit structures. Today's vehicle software architecture is still very much based on functions and therefore does not enable reuse of software components. Software architecture needs to become independent

of function architecture, decoupled by adequate and high-level software interfaces and decomposed in fine grained, reusable building blocks of software.

4 Acting within the Vehicle's Context

After evaluating the vehicle's surroundings the derived driving decision needs to be carried out. In [WZKT10] a semi-reactive approach for generating trajectories is outlined. Its input are abstract commands from the previous layer, which are used to calculate a desired trajectory with respect to long-term goals like keeping a desired velocity while following other vehicles, and to short-term goals like avoiding collisions.

In [LF09] an approach is outlined to generate dynamically feasible trajectories with the goal to travel on high speed. It is based on a lattice state space running with real-time performance by using a combination of high-resolution action space around the vehicle and a low-resolution action space elsewhere. In contrast, [ZS09] presents a method that reduces the required amount of nodes to model all possible vehicle's motions. An approach to stabilize state trajectories for inner-city speeds up to $6 \frac{m}{s}$ is outlined in [WGB10].

In [KPJ⁺10], a system is described, which allows to control precisely an autonomously driving vehicle to slide into a sideways parking spot. Although there is only a limited practical usage for the daily usage in urban environments, the authors showed that they were able to control the vehicle repeatedly even in this extremely dynamical maneuver.

From a software engineering point of view, modeling of control algorithms is very mature. Challenges are the reliability even in degrading situations, in particular as we cannot reliably predict the absent driver (e.g. due cognitive distraction) to take over within seconds as a fallback. Low-level safety and reliability is the key here. As a second challenge, we face that adaptivity of higher functions to customer specific needs enforce capabilities of dynamic updates and enhancements of car software. Among others this imposes an additional security problem. Thus, we need an appropriate software architecture that allows to aggregate and cummlate information from real and virtual sensors, but also produces efficient shortcuts from sensors to actuators in case of emergencies.

5 Monitoring & Evaluating the Vehicle's Performance

As shown in Fig. 1(b) the development of autonomously driving vehicles is also supported by visualizing and analyzing recorded data from various test drives to comprehend a driving decision. However, this data can only be analyzed when all required sensors are mounted and calibrated accordingly. But for testing the first sketches of algorithms anyhow, existing data sets from the DARPA Urban Challenge [HAO⁺10] or from Ford [PME11] are available to validate the correctness of data handling and processing. Furthermore, the Stanford Track Collection [TLT11] also provides nearly 14,000 labeled tracks for further inspection. These tracks are recorded using a Velodyne HDL-64E S2 LIDAR

for various street scenarios. For example, this database could be used to develop algorithms that were able to classify traffic participants like pedestrians and bicyclists to evaluate the reliability of another detecting sensor like a vision system.

From a software engineering point of view, this availability of data provides an interesting opportunity to define an enhanced process to develop software. New versions of software can be tested in virtual environments completely detached from hardware. This enables software to drive many millions virtual miles much faster than any “real time” HiL could allow. Simulation is the key for software quality management and software and function architectures can be defined in such a way that various combination of software components and their interplay can be tested while its context is being simulated through appropriate mocks.

6 Engineering the Autonomously Driving Vehicle’s Software

The previous sections summarized the research results accounting aspects, which were beyond the requirements of the DARPA Urban Challenge according to the general system architecture. However from a software engineer’s point of view, aspects like the development process itself, integration and quality assurance for components or the entire system, or the re-use of existing components come into mind. However, these aspects have hardly emerged from finalist teams within the last five years. Only the authors of [BDWL11] address the actual reuse of software components that is related to software engineering.

However, concerning the increasing complexity and uncertainty of sensed input data, real test drives as carried out to find optimal parameters for the 2007 competition are insufficient. Instead, a thorough engineering approach to develop the software of these systems is required, which supports the development already at early stages by virtual environments. Regarding the tight schedule for the 2007 DARPA Urban Challenge together with the limited vehicles’ surroundings, the main focus of all teams was on the actual performance during the qualification and the final. However, modern comfort and active safety systems has to operate properly even in hardly foreseeable traffic situations. Thus, today’s software engineering must be extended by approaches, which explicitly embrace simulation environments especially for sensor-based data.

In [Ber10] and [?], foundations, modeling aspects, and applications for such a virtual environment are outlined and discussed. One of the main ideas is to model the vehicle’s surroundings in terms of a single point of truth. From this hierarchical environmental model, particular aspects for the relevant data processing layers (cf. Fig. 1(b)) are derived and used within a dedicated simulation context: Its object-oriented representation is used to serve the decision layer, while the annotated various 3D models are used to generate the required raw sensor data for the perception layer. During the development of “Caroline” and a succeeding project at the University of California, Berkeley, selected aspects from this simulation-based software engineering approach were evaluated.

7 Conclusion and Outlook

In this paper, contributions from finalist teams of the 2007 DARPA Urban Challenge from the last five years were summarized with a focus on aspects, which were explicitly excluded from that competition. Regarding the goal to develop an anticipatorily driving vehicle that is reliable enough for the daily use, current research is on a solid way. This is also documented by an increasing number of advanced driver assistance and safety systems, which increase the safety not only for the occupants but also for other traffic participants like pedestrians. In the upcoming future, we will encounter more and more systems that base on a reliable environmental model and which will be increasingly interconnected to serve their users.

These assisting and interconnected sensor- and actor-based systems are currently classified as *Cyber-Physical Systems*, which enable newly arising services and business models to alleviate our daily life [GB12]. However, the engineering of these increasingly intelligent, interconnected, and autonomously acting systems demands for more elaborated and partially new techniques and methods [GRSS12]; also formal methods can be successfully applied to focus on the *correct* implementation of the *right* requirements [SHGB11].

For the long-term success of these systems, it is fundamentally important to not only regard the desired behavior and functionality of these systems but even more to analyze and model *timing properties* as part of their semantic correctness within their real domain [Lee12]. Without these adapted or newly developed methods it will hardly be possible to evaluate the correct behavior of a system within its application domain. Furthermore in future, the evaluation of a system in reality must be extended by virtual approaches to ensure a valid behavior in critical or unforeseeable situations. This trend can already be seen by the research's results in the evolution of autonomously driving vehicles.

References

- [BCG⁺09] Alberto Broggi, Pietro Cerri, Stefano Ghidoni, Paolo Grisleri, and Ho Gi Jung. A New Approach to Urban Pedestrian Detection for Automatic Braking. *IEEE Transactions on Intelligent Transportation Systems*, 10(4):594–605, December 2009.
- [BDWL11] Christopher R. Baker, John M. Dolan, Shige Wang, and Bakhtiar B. Litkouhi. Toward adaptation and reuse of advanced robotic software. In *Proceedings of the IEEE International Conference on Robotics and Automation*, pages 6071–6077, Shanghai, China, May 2011.
- [Ber10] Christian Berger. *Automating Acceptance Tests for Sensor- and Actuator-based Systems on the Example of Autonomous Vehicles*. Shaker Verlag, Aachener Informatik-Berichte, Software Engineering Band 6, Aachen, Germany, 2010.
- [CRZ11] Hyunggi Cho, Paul E. Rybski, and Wende Zhang. Vision-based 3D Bicycle Tracking using Deformable Part Model and Interacting Multiple Model Filter. In *Proceedings of the IEEE International Conference on Robotics and Automation*, pages 4391–4398, Shanghai, China, May 2011.

- [CWW11] Paul Czerwionka, Miao Wang, and Fabian Wiesel. Optimized Route Network Graph as Map Reference for Autonomous Cars Operating on German Autobahn. In *Proceedings of the 5th International Conference on Automation, Robotics and Applications*, pages 78–83, Wellington, New Zealand, December 2011.
- [DTMD09] Dmitri Dolgov, Sebastian Thrun, Michael Montemerlo, and James Diebel. Path Planning for Autonomous Driving in Unknown Environments. In *Springer Tracts in Advanced Robotics*, pages 55–64, Berlin Heidelberg, 2009. Springer.
- [FU11] Nathaniel Fairfield and Chris Urmson. Traffic Light Mapping and Detection. In *Proceedings of the IEEE International Conference on Robotics and Automation*, pages 5421–5426, Shanghai, China, May 2011.
- [GB12] Eva Geisberger and Manfred Broy, editors. *agendaCPS - Integrierte Forschungsagenda Cyber-Physical Systems (acatech STUDIE)*. Springer Verlag, Heidelberg, 2012.
- [GRSS12] Holger Giese, Bernhard Rumpe, Bernhard Schätz, and Janos Sztipanovits. Science and Engineering of Cyber-Physical Systems. *Dagstuhl Reports*, 1(11):1–22, 2012.
- [GSL11] Haifeng Gong, Jack Sim, Maxim Likhachev, and Jianbo Shi. Multi-hypothesis Motion Planning for Visual Object Tracking. In *Proceedings of the International Conference on Computer Vision*, pages 619–626, Barcelona, Spain, November 2011.
- [HAO⁺10] Albert S. Huang, Matthew E. Antone, Edwin Olson, Luke Fletcher, David Moore, Seth J. Teller, and John J. Leonard. A High-rate, Heterogeneous Data Set From The DARPA Urban Challenge. *The International Journal of Robotics Research*, 29(13):1595–1601, November 2010.
- [HC11] Jason Hardy and Mark Campbell. Clustering Obstacle Predictions to Improve Contingency Planning for Autonomous Road Vehicles in Congested Environments. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 1605–1611, San Francisco, CA, USA, September 2011. IEEE.
- [HW11] Kai Homeier and Lars Wolf. RoadGraph : High level sensor data fusion between objects and street network. In *Proceedings of the International IEEE Conference on Intelligent Transportation Systems*, pages 1380–1385, Washington, DC, USA, October 2011.
- [JHMBH10] Prasanth Jeevan, Frank Harchut, Bernhard Mueller-Bessler, and Burkhard Huhnke. Realizing Autonomous Valet Parking with Automotive Grade Sensors. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 3824–3829, Taipei, Taiwan, October 2010.
- [KPJ⁺10] J. Zico Kolter, Christian Plagemann, David T. Jackson, Andrew Y. Ng, and Sebastian Thrun. A Probabilistic Approach to Mixed Open-loop and Closed-loop Control, with Application to Extreme Autonomous Driving. In *Proceedings of the IEEE International Conference on Robotics and Automation*, pages 839–845, Anchorage, AL, USA, May 2010.
- [LADT11] Jesse Levinson, Jake Askeland, Jennifer Dolson, and Sebastian Thrun. Traffic Light Mapping, Localization, and State Detection for Autonomous Vehicles. In *Proceedings of the IEEE International Conference on Robotics and Automation*, pages 5784–5791, Shanghai, China, May 2011.
- [Lee12] Edward A Lee. Time for High-Confidence Cyber-Physical Systems. In *Proceedings of the Performance Metrics for Intelligent Systems Workshop*, pages 1–30, College Park, MD, USA, March 2012.

- [LF09] Maxim Likhachev and Dave Ferguson. Planning Long Dynamically Feasible Maneuvers for Autonomous Vehicles. *International Journal of Robotics Research*, 28(8):933–945, August 2009.
- [LMT07] Jesse Levinson, Michael Montemerlo, and Sebastian Thrun. Map-Based Precision Vehicle Localization in Urban Environments. In *Robotics: Science and Systems*, 2007.
- [LT10] Jesse Levinson and Sebastian Thrun. Robust Vehicle Localization in Urban Environments Using Probabilistic Maps. In *Proceedings of the IEEE International Conference on Robotics and Automation*, pages 4372–4378, Anchorage, AL, USA, May 2010.
- [Nan11] Hao Nan. Car takes long drive - by itself, August 2011.
- [NHF⁺11] Tobias Nothdurft, Peter Hecker, Tobias Frankiewicz, Jan Gacnik, and Frank Köster. Reliable Information Aggregation and Exchange for Autonomous Vehicles. In *Proceedings of the IEEE Vehicular Technology Conference*, pages 1–5, San Francisco, CA, USA, September 2011.
- [NHO⁺11] Tobias Nothdurft, Peter Hecker, Sebastian Ohl, Falko Saust, Markus Maurer, Andreas Reschka, and Jürgen Rüdiger Böhmer. Stadtpilot: First Fully Autonomous Test Drives in Urban Traffic. In *Proceedings of the International IEEE Conference on Intelligent Transportation Systems*, pages 919–924, Washington, DC, USA, October 2011.
- [PME11] Gaurav Pandey, James R McBride, and Ryan M Eustice. Ford Campus vision and lidar data set. *The International Journal of Robotics Research*, 30(13):1543–1552, March 2011.
- [RBL⁺08] Fred W. Rauskolb, Kai Berger, Christian Lipski, Marcus Magnor, Karsten Cornelien, Jan Effertz, Thomas Form, Fabian Graefe, Sebastian Ohl, Walter Schumacher, Jörn-Marten Wille, Peter Hecker, Tobias Nothdurft, Michael Doering, Kai Homeier, Johannes Morgenroth, Lars Wolf, Christian Basarke, Christian Berger, Tim Gülke, Felix Klose, and Bernhard Rumpe. Caroline: An Autonomously Driving Vehicle for Urban Environments. *Journal of Field Robotics*, 25(9):674–724, September 2008.
- [SFG⁺10] Todd Sullivan, Richard Frankel, Olafur Gudmundsson, Brett Miller, Jordan Potter, Salik Syed, Doreen Hoang, Jae min John, Ki-Shui Liao, Pasha Nahass, Amanda Schwab, and Jessica Yuan. Assisted Highway Lane Changing with RASCL. In *Proceedings of the AAAI Spring Symposium: Embedded Reasoning*, pages 1–8, 2010.
- [SHGB11] Sebastian Siegl, Kai-Steffen Hielscher, Reinhard German, and Christian Berger. Automated Testing of Embedded Automotive Systems from Requirement Specification Models. In *Proceedings of the 12th IEEE Latin-American Test Workshop*, pages 1–6, Porto de Galinhas, Brazil, March 2011.
- [SP12] Christoph Stiller and Oliver Pink. Real-Time Knowledge for Cooperative Cognitive Automobiles. In Samarjit Chakraborty and Jörg Eberspächer, editors, *Advances in Real-Time Systems*, pages 341–355, Heidelberg Dordrecht London New York, 2012. Springer.
- [SU08] Young-Woo Seo and Chris Urmson. A Perception Mechanism for Supporting Autonomous Intersection Handling in Urban Driving. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 1830–1835, Nice, France, September 2008.
- [SU09] Young-Woo Seo and Chris Urmson. Utilizing Prior Information to Enhance Self-Supervised Aerial Image Analysis for Extracting Parking Lot Structures. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 339–344, St. Louis, MO, USA, October 2009.

- [SUWL10] Young-Woo Seo, Chris Urmson, David Wettergreen, and Jin-Woo Lee. Building Lane-Graphs for Autonomous Parking. In *Proceedings of IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS-10)*, pages 6052–6057, Taipei, Taiwan, October 2010.
- [Thr10a] Sebastian Thrun. Toward Robotic Cars. *Communications of the ACM*, 53(4):99–106, April 2010.
- [Thr10b] Sebastian Thrun. What we’re driving at, April 2010.
- [TLT11] Alex Teichman, Jesse Levinson, and Sebastian Thrun. Towards 3D Object Recognition via Classification of Arbitrary Object Tracks. In *Proceedings of the IEEE International Conference on Robotics and Automation*, pages 4034–4041, Shanghai, China, May 2011.
- [VGZD07] Stefan Vacek, Tobias Gindele, Johann Marius Zöllner, and Rüdiger Dillmann. Using case-based reasoning for autonomous vehicle guidance. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 4271–4276, San Diego, CA, USA, October 2007.
- [WDSL11] Junqing Wei, John M. Dolan, Jarrod M. Snider, and Bakhtiar Litkouhi. A Point-based MDP for Robust Single-Lane Autonomous Driving Behavior under Uncertainties. In *Proceedings of the IEEE International Conference on Robotics and Automation*, pages 2586–2592, Shanghai, China, May 2011.
- [WGB10] Moritz Werling, Lutz Gröll, and Georg Bretthauer. Invariant Trajectory Tracking With a Full-Size Autonomous Road Vehicle. *IEEE Transactions on Robotics*, 26(4):758–765, August 2010.
- [WGR11a] Miao Wang, Tinosch Ganjineh, and Raúl Rojas. Action Annotated Trajectory Generation for Autonomous Maneuvers on Structured Road Networks. In *Proceedings of the 5th International Conference on Automation, Robotics and Applications*, pages 67–72, Wellington, New Zealand, December 2011.
- [WGR11b] Jianxin Wu, Christopher Geyer, and James M. Rehg. Real-Time Human Detection Using Contour Cues. In *Proceedings of the IEEE International Conference on Robotics and Automation*, pages 860–867, Shanghai, China, May 2011.
- [WZKT10] Moritz Werling, Julius Ziegler, Sören Kammel, and Sebastian Thrun. Optimal Trajectory Generation for Dynamic Street Scenarios in a Frenét Frame. In *Proceedings of the IEEE International Conference on Robotics and Automation*, pages 987–993, Anchorage, AL, USA, May 2010.
- [ZRG⁺09] Brian D. Ziebart, Nathan D. Ratliff, Garratt Gallagher, Christoph Mertz, Kevin M. Peterson, James Andrew Bagnell, Martial Hebert, Anind K. Dey, and Siddhartha S. Srinivasa. Planning-based Prediction for Pedestrians. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 3931–3936, St. Louis, MO, USA, October 2009.
- [ZS09] Julius Ziegler and Christoph Stiller. Spatiotemporal state lattices for fast trajectory planning in dynamic on-road driving scenarios. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 1879–1884, St. Louis, MO, USA, October 2009.