Report of Dagstuhl Seminar 15462

The Mobile Revolution – Machine Intelligence for Autonomous Vehicles

Edited by

Wolfram Burgard[1], Uwe Franke[2], Markus Enzweiler[3], and Mohan Trivedi[4]

1 **Universität Freiburg, DE,** burgard@informatik.uni-freiburg.de
2 **Daimler AG – Sindelfingen, DE,** uwe.franke@daimler.com
3 **Daimler AG – Böblingen, DE,** markus.enzweiler@daimler.com
4 **University of California, San Diego – La Jolla, US,** mtrivedi@ucsd.edu

Abstract

This report documents the Dagstuhl Seminar 15462 "The Mobile Revolution – Machine Intelligence for Autonomous Vehicles". The seminar has discussed the state-of-the-art and provided a consistent vision on the topic of intelligent autonomous vehicles. It has served as a communication platform between the various sub-communities involved, by bringing together key persons in industry and academia in their respective fields. Additionally, relations between different disciplines of intelligent transportation systems have be identified and exploited. The seminar has allowed its participants to bridge the gap between foundational research and real-world applications by identifying further research directions and initiating interdisciplinary collaborations.

Seminar November 10–13, 2015 – http://www.dagstuhl.de/15462
1998 ACM Subject Classification Autonomous vehicles, Robotics, Vision and Scene Understanding
Keywords and phrases autonomous driving, digital maps, environment perception, machine intelligence, robotics, situation awareness, traffic safety
Digital Object Identifier 10.4230/DagRep.5.11.62

1 Executive Summary

Markus Enzweiler

Motivation and Perspective

Machine intelligence, robotics and computer vision, formerly rather peripheral disciplines of computer science, are in fact already with us today and have a familiar embodiment – the modern vehicle. Systems that are currently available strongly couple interdisciplinary fundamental research with complex practical realizations. The vision of autonomous vehicles in particular has a surprisingly long history with first prototypical implementations going back to the early 1980s. What started then as a dream of pioneers such as Ernst Dickmanns is actually happening right now – we are on the verge of a mobile revolution with self-driving vehicles as its central foundation. The tremendous progress made in the last years has been sparked by the increased methodical and technically availability of better sensors, sophisticated algorithms, faster computers and more data.

But, we are not quite there yet. Autonomous systems make extreme demands on system performance, quality, availability, reliability and verification that significantly increase with the rising degree of automation. Such diverse requirements give rise to numerous problems and open questions that are currently addressed in substantial academic and industrial research activities in many fields of computer science and engineering. Extraordinarily positive innovation effects result from the knowledge transfer between industry and academia, as successfully demonstrated by initiatives such as Uni-DAS or DRIVE-U. The increasing relevance and interest in the computer science community, particularly in the fields of robotics, computer vision and machine learning is evident through an abundance of papers and workshops at major computer science conferences.

This seminar has brought together the leading experts from both academia and industry to discuss the state-of-the-art, identify further research directions and refine the overall vision of intelligent autonomous vehicles into a consistent and practicable picture.

Seminar Topics and Structure

The ultimate design goal for autonomous systems is to mimic human behavior in terms of understanding and effortlessly acting within a dynamic human-inhabited environment. Although artificial sensors emulating the human sensory systems are nowadays widely available, current autonomous systems are still far behind humans in terms of understanding and acting in real-world environments. The chief reason is the (theoretical and practical) unavailability of methods to reliably perform perception, recognition, understanding and action on a broad scale, i.e. not limited to isolated problems.

Following the classical perception-action cycle, the central topics of the seminar have evolved around four key questions posed from the perspective of an autonomous vehicle:
- What do I perceive and how can I interpret this?
- Where am I and what do I do next?
- How can I build up experience and learn?
- Am I capable of this task?

More specifically, the seminar has stimulated research and discussions through several talks on the following topics:
- Intelligent Robotics
- Digital Maps
- Human-centered Intelligent Vehicles
- Verification and Validation
- Limitations and Perspectives

The seminar was held in a very interactive workshop style allowing for ample time for thorough discussions. There were four main sessions with talks and discussions, c.f. the seminar schedule in Section 4, focusing on autonomous driving projects, mapping and localization, sensing, as well as evaluation and approval. The first session on state-of-the-art autonomous driving projects has been co-organized with Seminar 15461 as a joint session.

2 Table of Contents

3 Overview of Talks

3.1 Realizing Self-Driving Cars

Andreas Wendel (Google X, Mountain View, US)

Self-driving vehicles are coming. They will save lives, save time and offer mobility to those who otherwise don't have it. Eventually they will reshape the world we live in. A dedicated team at Google has spent the last few years moving self-driving vehicles closer to reality. New algorithms, increased processing power, innovative sensors and massive amounts of data enable our vehicles to see further, understand more and handle a wide variety of challenging driving scenarios. Our vehicles have driven over a million miles on highways, suburban and urban streets. Through this journey, we've learned a lot; not just about how to drive, but about interacting with drivers, users and others on the road, and about what it takes to bring an incredibly complex system to fruition. In my talk, I share some insights in how the technology works, how we have rolled out our new prototype vehicles to public roads, and which edge case situations we have to solve.

3.2 Autonomous Vehicles – Relations to Human Intelligence

Klaus Bengler (TU München, DE)

Human behavior and change of human behavior is basis not only for critical incidents and accidents but also for efficient traffic and mitigation. The introduction of automation is going to replace or even should support these mechanisms. HMI design plays a dominant role here. Furthermore the introduction of automation will lead to behavioral changes regarding users of the automation and other traffic participants. These effects need further research.

3.3 The Use of 3D Prior Maps in Automated Driving

Ryan Eustice (University of Michigan – Ann Arbor, US)

Self-driving test vehicles have become a reality on roadways and there is an ever present push toward making them a consumer product in the not so distant future. In this talk, I will give an overview of some of our on-going work (in collaboration with Ford Motor Company) in full-scale automated driving. In particular, we'll look at some of our successes in high definition map building and precision localization, including our recent work in cross-modality localization using vision within a priori LIDAR maps. We'll also reflect upon the challenges ahead in perception and human factors.

3.4 HERE Vision of the Future for Autonomous Vehicles

D. Scott Williamson (HERE – Chicago, US) and Alex Goldberg (HERE – Carlsbad, US)

At HERE we are trying to answer three questions: "Where exactly am I?", "What Lies Ahead?", and "How can I get there comfortably?". Maps are needed for autonomous driving as an extended sensor of the road ahead in order to make autonomous driving safer and more comfortable. Comfort is enabled by permitting planning beyond sensor range, and providing information about historical driving behaviour. We present our vision and strategies to address these challenges. Our vision includes a map with accurate lane level geometry, localization features, live traffic events, and speed profiles. This information is delivered via a cloud service that is always enabled providing the freshest data available. In order to achieve the freshness required for autonomous driving we see great potential in utilizing sensor feedback from autonomous vehicles. We see a need for industry alignment on features required for localization, this requires active area of industry collaboration.

3.5 How to Address the Approval Trap for Autonomous Vehicles

Hermann Winner (TU Darmstadt, DE), Maren Graupner (TU Darmstadt, DE), and Walther Wachenfeld (TU Darmstadt, DE)

Autonomous vehicles will pass the technology readiness level of prototype demonstrators in an operational environment soon. Thereby the human ability to control the vehicle must be fully replaced by a technical system. Such a cognitive system that perceives and processes the complex world in public traffic has never before been approved for series production. The so called approval trap appears which means that a ready to use developed autonomous vehicle cannot be released due to the lack of safety validation concepts. What could be the way out? Two approaches with different introduction strategies will lead to gain the needed knowledge for validation. For both efficient test tools and systematical reduction of test cases have to be developed. A method of decomposition of the entire automation process into functional layer corresponding to human driving skills is introduced. This decomposition opens an orthogonal way for testing machine driving efficiently. But, in any case the safety validation has to be part of the first introduction of autonomous driving by risk limitation.

3.6 Driving in Unstructured Environments

Hans-Joachim Wünsche (Universität der Bundeswehr – München, DE)

I started my presentation with my definition of "unstructured environment": hard to "see´´ road/paths (view through on-board camera onto a muddy forest road during Elrob 2007), poor maps and very poor GPS (from Elrob 2009). This motivates my research: concentrate on perception (even it is very tough), forget precise metric maps (rather create topological maps containing relevant, and observable /re-cognizable landmarks and try to get by without GPS. Examples of perception shown in several videos include detection of forest roads and unknown off-road intersections by vision and by raw data and feature based fusion of vision, lo-light, thermal cameras and Lidar data, both during day time and data in the middle of the night. Raw data fusion is also used to build rich, colored 3D-terrain maps of the immediate vehicle surrounding from Lidar point clouds colored through a color camera yawing left & right and encoding terrain steepness and info (if available) from probably available maps such as vegetation. Landmarks are perceived and have to be so compact as to transmit them via a 9600 baud radio link to other vehicles, which then use this info to build topological maps. These help to navigate autonomously in order to follow each other at larger distances, or to go from A to B without GPS, just localizing relative to landmarks. A video showed successful action, even if the landmarks are seen "from the rear" because the path is driven in reverse order than recorded. The talk closed with emphasizing human like behavior also for "object relational" trajectory planning and recursive situation assessment, to concentrate on what's relevant and less on what's available.

3.7 Human Factors in Intelligent Vehicles

Mohan Trivedi (University of California, San Diego – La Jolla, US)

Designing fully autonomous robotic vehicles which can drive on roads does not require models of drivers and how they interact with vehicles. In contrast – design of humanized "intelligent" vehicles especially those for active safety that prevent accidents, requires understanding of human behavior, modeling of human-vehicle interaction, activities inside the vehicle, and prediction of human intent. We present an overview of a fifteen year old journey of research into such activities and share our findings and experience of basic research as well as road-tested experimental studies in the real-world driving conditions. A holistic framework that utilizes all contextual cues from looking outside as well as inside the vehicle and a sparse Bayesian learning framework has provided convincing results for deployment of predictive, robust and reliable performances.

3.8 Towards Cooperative Autonomous Vehicles

Christoph Stiller (KIT – Karlsruher Institut für Technologie, DE)

In my talk I first revisited the challenges in Automated Driving and outlined the lessons learned from it. While the Urban Challenge 2007 showed that map-based driving with Laser- Scanners and DGPS/INS was feasible in urban areas with low buildings, the Grand Cooperative Driving Challenge in the Netherlands 2011 demonstrated cooperative maneuvers between cars fro, teams that have developed their systems independently for the first time. In collaboration of FZI/KIT with Daimler, the Bertha Benz Route automation showed that autonomous vehicles mainly running with vision sensors may navigate autonomously through a densely populated area in Germany. The talk closed with open issues and remaining challenges that need to be addressed to make level 4/5 automated vehicles a reality.

3.9 Sensors and Perceptual Algorithms

Michael James (Toyota Research Institute North America- Ann Arbor, US)

Highly automated vehicles depend on a wide range of sensors, information sources, and perceptual algorithms to make sense of their surroundings. In our view, this includes direct sensing technologies such as cameras, lidar, and radar; sources of information such as high-definition maps; and sources of remote data such as V2X. We present an overview of how these sensors are used at Toyota Research Institute North America in our highly automated test vehicle, similar in many respects to other research vehicles around the world. Then we present an overview of current and in-development sensors and their capabilities and limitations, as well as a set of potential open questions about how sensors and perception algorithms can be developed in the future.

3.10 Toward Fully Automated Driving

Jan Becker (BOSCH Research Center – Palo Alto, US)

In this talk, we present Bosch's approach to highly and fully automated driving. We define our vision and roadmap and introduce highway pilot as an example for a highly automated system. Then we present underlying technologies, localization and mapping, perception and planning. The talk emphasizes remaining challenges from the current state of the art to series production. Notable challenges are surround sensing robust in all use cases, safety and security i.e. protection against technical failures as well as against deliberate attacks, global standards for legislation and new liability regulations, commercially available precise and up-to-date maps, and new system architecture requirements which includes redundancies for sensing, ECUs, and actuators. We close with an analysis of human factor requirements, specifically regarding transition between automated and manual driving, as well as of the user experience on a highly or fully automated vehicle.

3.11 Coast to Coast and Urban Driving Experience

Serge Lambermont (Delphi Labs – Mountain View, US)

This talk outlines the market development and market drivers for autonomous vehicles. It briefly describes the conventional automotive market with increased ADAS functionality and level 2 traffic jam assist as well as highway pilot. After this adjacent markets as Mobility on Demand are described. Subsequently the sensor impact and new sensor technologies are described. Then, experiences on an automated drive from San Francisco to New York City, as well as urban drives in Las Vegas and Silicon Valley are provided. In the last part, situations and obstacle are described in an occurrence versus severity plane.

3.12 Autonomous Automobiles – Current Challenges in Research and Development

Markus Maurer (TU Braunschweig, DE)

In this talk I summarized the current challenges in the research and development of autonomous automobiles. To make sure everybody had the same understanding of "Autonomous Driving" I started with a short definition of autonomy. I also made my personal perspective transparent: it is basically a systems perspective. In the main part of the talk I addressed the following research fields on the way towards autonomous driving: requirements on infrastructure, system architecture, security, functional safety, perception, decision making, cooperative testing, open systems, acceptance and risk management and education.

4 Schedule

Wednesday, November 11, 2015 (partly shared with seminar 15461)

09:00–09:15	Welcome Address	Organizers
09:15–10:35	Autonomous Driving Projects (I)	Christoph Stiller, Raul Rojas
10:35–10:55	Coffee Break	
10:55–12:15	Autonomous Driving Projects (II)	Markus Maurer, Andreas Wendel
12:15–13:30	Lunch	
13:30–14:30	Introduction Round	All Participants
14:30–15:50	Autonomous Driving Projects (III)	Klaus Dietmayer, Wolfram Burgard
15:50–16:10	Coffee Break	
16:10–17:30	Autonomous Driving Projects (IV)	Hans-Joachim Wünsche, Jan Becker

Thursday, November 12, 2015

09:00–10:00	Session: Mapping and Localization (Chair: Mohan Trivedi)	Alex Goldberg, Scott Williamson, Ryan Eustice
10:00–10:10	Coffee Break	
10:10–12:00	Group Work on Mapping and Localization	
12:00–13:30	Lunch	
13:30–14:30	Session: Sensing (Chair: Uwe Franke)	Serge Lambermont, Michael James
14:30–14:40	Coffee Break	
14:40–17:00	Group Work on Sensing	

Friday, November 13, 2015

09:00–10:00	Session: Evaluation and Approval (Chair: Markus Enzweiler)	Hermann Winner
10:00–10:20	Coffee Break	
10:20–11:40	Session: Humans in Autonomous Cars (Chair: Markus Enzweiler)	Mohan Trivedi, Klaus Bengler
11:40–12:00	Workshop Closing	
12:00–13:30	Lunch	

Participants

- Michael Aeberhard
BMW AG – München, DE
- Jan Becker
BOSCH Research Center – Palo Alto, US
- Klaus Bengler
TU München, DE
- Claus Brenner
Leibniz Univ. Hannover, DE
- Wolfram Burgard
Universität Freiburg, DE
- Erik Coelingh
Volvo Car Corporation – Göteborg, SE
- Michael Darms
Volkswagen AG – Wolfsburg, DE
- Markus Enzweiler
Daimler AG – Böblingen, DE
- Ryan Eustice
University of Michigan – Ann Arbor, US
- Uwe Franke
Daimler AG – Sindelfingen, DE
- Dariu M. Gavrila
Daimler R&D – Ulm, DE
- Alex Goldberg
HERE – Carlsbad, US
- Ralf G. Herrtwich
Daimler AG – Böblingen, DE
- Ulrich Hofmann
Audi AG – Ingolstadt, DE
- Michael James
Toyota Research Institute North America – Ann Arbor, US
- Serge Lambermont
Delphi Labs – Mountain View, US
- Antonio M. Lápez Pena
Autonomus University of Barcelona, ES
- Chris Mansley
BOSCH Research Center – Palo Alto, US
- Markus Maurer
TU Braunschweig, DE
- Karsten Mühlmann
Robert Bosch GmbH – Heilbronn, DE
- Urs Muller
NVIDIA Corp. – Morganville, US
- Michel Parent
INRIA – Le Chesnay, FR
- Mikael Persson
Linköping University, SE
- Raul Rojas
FU Berlin, DE
- Torsten Sattler
ETH Zürich, CH
- Steven E. Shladover
Univ. of California, Berkeley, US
- Christoph Stiller
KIT – Karlsruher Institut für Technologie, DE
- Matthias Strauß
Continental Teves AG – Frankfurt, DE
- Mohan Trivedi
University of California, San Diego – La Jolla, US
- Sadayuki Tsugawa
AIST – Ibaraki, JP
- D. Scott Williamson
HERE – Chicago, US
- Hermann Winner
TU Darmstadt, DE
- Hans-Joachim Wünsche
Universität der Bundeswehr – München, DE