

In Cyber Valley, researchers are working on solutions that make autonomous driving easily scalable.

End-to-End Training

802

Professor Dr. Andreas Geiger is Head of the Autonomous Vision research group at the Max Planck Institute for Intelligent Systems (MPI-IS) and Professor for learning-based computer vision and autonomous vision at the University of Tübingen. In this interview, he talks about the challenges of developing self-driving cars, explains how attractive German universities are by international comparison and elaborates on what has to be done to keep young talents in the country.

In 2018, you received the IEEE PAMI Young Researcher Award for your outstanding contribution to bridging the gap between computer vision, machine learning, and robotics. Can you briefly describe what the award means to you?

The award means a lot to me because it recognizes the international importance of my work and shows that we are on par with the best computer vision research labs in the world. I was the first German researcher to receive it and the third researcher in Europe. The award is the most prestigious distinction in the field of computer vision for a young researcher.

What exactly did the award honor?

I was awarded for my research on self-driving cars. I first developed algorithms and approaches to achieve scene understanding in my dissertation at the Karlsruhe Institute of Technology. Since I wanted to work with real data, I equipped a test vehicle with comprehensive sensor technology – several cameras, lidars, and GPS. At some point, we decided to make the tediously collected data available to the general public. This created a by-product of my dissertation: the KITTI Benchmark, which was created in 2012 and has become one of the most influential data sets in the field of autonomous driving. Today, the KITTI Benchmark is the state of the art in the field of computer vision for evaluating algorithms.

In your words, what is the difference between control engineering and machine learning?

The lines between machine learning and control engineering are blurred and a matter of perspective. For a control engineer, perception is peripheral, for computer scientists, control engineering is peripheral. Personally, I think the more daunting challenges for autonomous driving are perception and AI-supported decision making. Compared to the control technology for a humanoid robot with 50 actuators and tactile sensors, the control system of a vehicle is relatively simple. Basically, a car is controlled only by steering, accelerator, and brakes. In addition, the industry has been working on vehicle control for a long time and has gained an immense amount of know-how accordingly.

Would you ride in an autonomous vehicle today?

Why not? I would not mind being a passenger in a Level-4 vehicle if the opportunity arose – There is usually a service employee in the car who can intervene if necessary.

And when do you think the first autonomous vehicles without service employees will be on the road?

Many industry representatives had promised that this would be the case in 2021. By now, many have retracted their promise and are being more realistic. I do not expect Level-5

autonomous driving in the next ten years because fundamental questions in the field of artificial intelligence have not been answered. Whether driving at Level 4 is successful depends on the defined framework conditions. In specific areas, under specific weather conditions, this may be possible in the next few years – as shown by Waymo. I suppose that we will start out with remote operators and speed limits. Tesla is a pioneer in this field, but I would be surprised to see an autonomous Tesla vehicle with Level 5 functions on the market within the next five years.

What are the greatest obstacles?

At present, we count one traffic fatality per 100 million miles. This shows that we humans have mastered driving quite well. An autonomous vehicle is intended to make fewer mistakes than the human driver and at best be better by a factor of ten or 100. It must therefore be safe in a whole range of different situations: For example, the cars have to perceive their surroundings at night, in the rain, and when it snows. Even though cameras are still far from being as good as the human eye, we have made considerable progress in the field of sensor technology in recent years. Autonomous vehicles must then be able to master busy or blocked roads. They must also be able to deal with arbitrary pedestrian behavior, reflections, and unpredictable and rare events.

>>

We are working on making simulation more efficient, because I am convinced that in the future, simulations that are as realistic as possible will become increasingly important for validation and training.

In order to train algorithms for these rare events, we therefore need an incredible amount of data. Another obstacle is that algorithms cannot perform causal inference, meaning they are not able to draw conclusions. Therefore, a high level of manual reprogramming is required in the systems. Add to this ethical and legal issues that must be clarified. As you can see, there is still some work to be done.

On which areas do you focus?

Our research group focuses on classical computer vision topics. For example, we are investigating how to improve depth perception and make it more robust. We also take a look at how algorithms can learn with less data. And we are working on making simulation more efficient, because I am convinced that in the future, simulations that are as realistic as possible will become increasingly important for validation and training. Lastly, we train algorithms for autonomous driving. In contrast to the automotive industry, which today works according to the classic modular approach, we are pursuing the approach of comprehensively trainable systems.

How does end-to-end training work and what are the benefits?

In end-to-end training, you try to view the whole system as one process, from perception to control, and represent it in one neural network. The system collects perception and control data of the vehicle, i.e., steering, acceleration and brake data. This gives us the advantage of training the system directly towards a goal instead of training individual modules towards subtasks, such as object recognition. We believe that these comprehensive models are the solution to better scale autonomous driving. Currently, these models are not as precise and robust as the modular concepts used in industry, for which a large number of engineers are working on individual modules. Once we manage data complexity, machine learning will allow us to move our system much faster into a new city, into new environments.

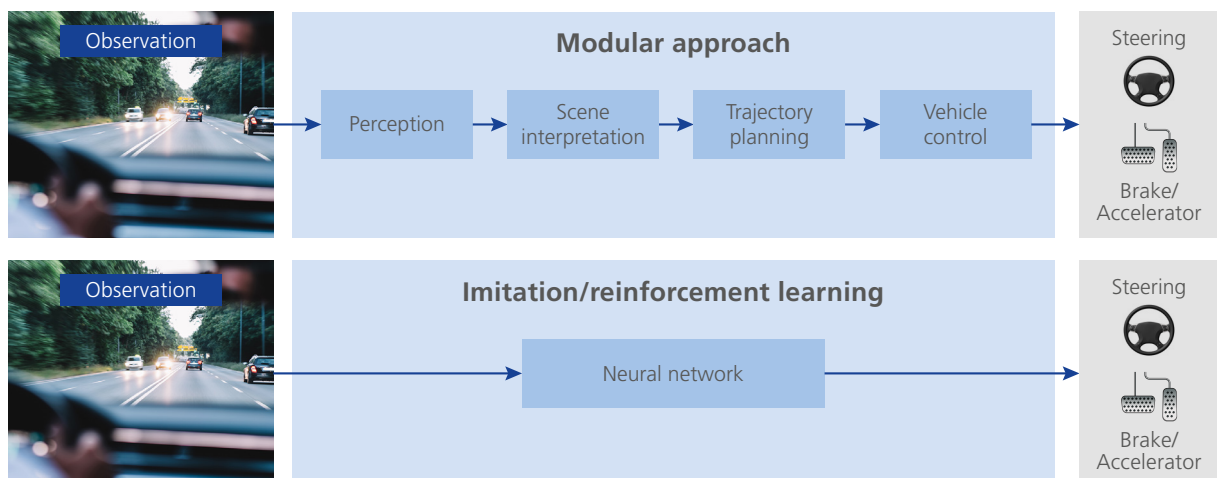
How closely do you work with the industry?

Even though the industry follows the modular approach, we cooperate in many subprojects with suppliers and

car manufacturers in the region. Our focus on the end-to-end approach is of great interest for the industry, even if they cannot apply it immediately. We are currently involved in the KI Delta Learning project, which analyzes self-learning methods for the automated processing of environment sensor data in the context of automated driving. The project has been commissioned by the Federal Ministry for Economic Affairs and Energy and involves leading industrial companies from the automotive industry as well as several universities, including the University of Tübingen.

What keeps you in the Cyber Valley in Tübingen, and how attractive are German universities by international comparison?

Europe is a strong player in academic research, and the automotive industry has a great interest in AI. The University of Tübingen and the Max Planck Institute are part of a large network of researchers who not only work on computer vision topics, but also apply AI in related disciplines, such as the neurosciences.



Picture credit: © A. Cardinale



Professor Dr. Andreas Geiger is head of the Autonomous Vision research group at the Max Planck Institute for Intelligent Systems (MPI-IS) and is Professor for learning-based computer vision and autonomous vision at the University of Tübingen.

KI Delta Learning

The aim of the KI Delta Learning research project is to evaluate the differences between domains and design new methods so that artificial intelligence can transfer existing knowledge from one domain to another and only has to learn the additional requirements, the specific “deltas”. This reduces the need for test data and accelerates the learning process when new knowledge has to be added.

In this network, we can learn from each other across disciplines. This makes working here very attractive. We also continue to network in various initiatives at the European level. One of these is ELLIS, the European Lab for Learning and Intelligent Systems, which promotes the exchange of information between institutes and doctoral students on machine learning and AI. You do not have to be in Silicon Valley to work with

the big companies there. Amazon is currently expanding its site here, Bosch is building a new location in our neighborhood, NVIDIA is sponsoring us, and I am working closely with Intel. However, we do have some catching up to do when it comes to start-ups.

What exactly do you mean?

One, we need to change the public mindset. Two, start-ups need more

support. The founder’s mindset is currently changing, but we need more incubators and less bureaucracy so that talented young people can turn their ideas into a reality here instead of being poached by the tech giants in the USA. Once they leave, they might never come back. Keeping talent here is vital.

Thank you for the interview.

The founder’s mindset is currently changing, but we need more incubators and less bureaucracy so that talented young people can turn their ideas into a reality here instead of being poached by the tech giants in the USA.