



# Rise of the Autonomous Machines

**Shaoshan Liu**, PerceptIn

**Jean-Luc Gaudiot**, University of California, Irvine

We are entering the age of autonomous machines, but many roadblocks exist on the path to make this a reality. We make a preliminary attempt at recognizing and categorizing the technical and nontechnical challenges of autonomous machines; for ten areas, we review current status, roadblocks, and potential research directions.

**A**s increasing numbers and kinds of autonomous machines enter our daily lives, the age of autonomous machines is upon us, and a whole new era of information technology begins. How did we get here? Before delving into autonomous machines, let us first review the evolution of computing machinery.

## THE AGE OF AUTONOMOUS MACHINES: THE SIXTH LAYER OF INFORMATION TECHNOLOGY

Information technology took off in the 1960s when Fairchild Semiconductor and Intel established its foundation by producing the first silicon microprocessors, with the attendant explosive growth of Silicon Valley. At first, although microprocessor technologies greatly improved industrial productivity, the general public had limited access to it. This changed in the 1980s with the advent of personal computers and, later, the Apple Macintosh and Microsoft Windows,

using a GUI. The second layer had been established, and the vision of ubiquitous computers, targeted for use by untrained personnel at home, had become a possibility.

With virtually everyone having access to computing power by the early 2000s, Yahoo and Google laid the third layer, connecting people—indirectly—with information available through search engines. With the third layer, the Internet, a core infrastructure to enable later layers, could become ubiquitous.

Beginning with Facebook in 2004, social networking sites created the fourth layer of information technology by allowing people to directly connect with each other, effectively moving the whole of human society to the World Wide Web. As the population of Internet-savvy people reached critical mass, the emergence of such applications as Airbnb (2008) and Uber (2009) was the basis for the fifth layer by providing direct Internet commerce services.

So far, each new layer of information technology, with its added refinements,

has incrementally improved popular access and demand. Note that, for most Internet commerce sites providing access to service providers through the Internet, it is still ultimately humans who are providing the services.

We have now entered the age of autonomous machines (the sixth layer), in which autonomous machines, such as service robots, autonomous drones, delivery robots, and autonomous vehicles, rather than humans, will provide services. Obviously, autonomous machines have the potential to completely upend our daily lives and our economy in the coming decade.

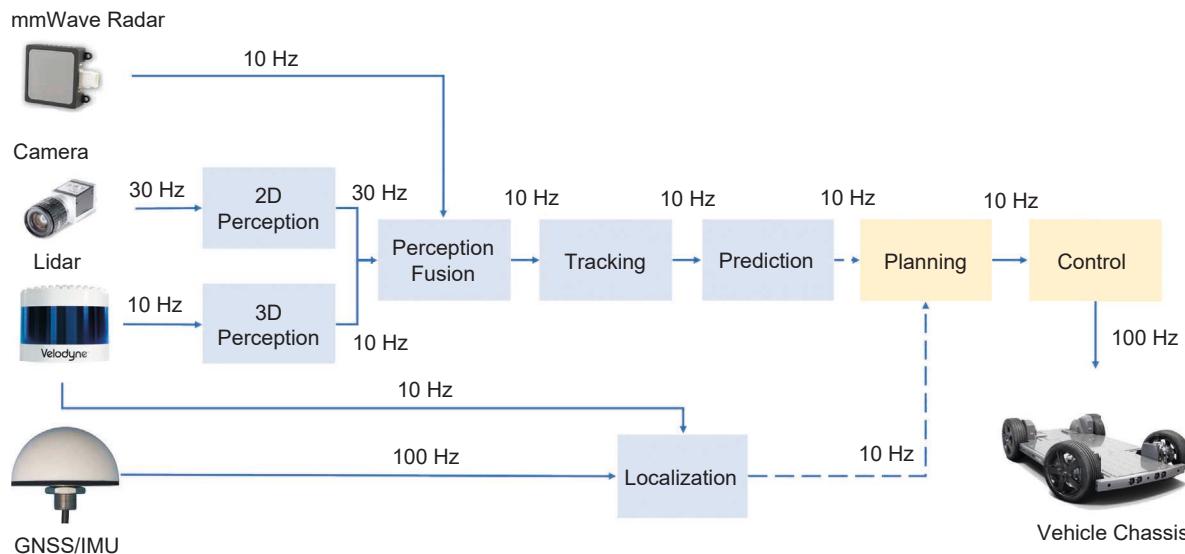
Yet, autonomous machines are extremely complex systems that integrate many segments of technology,<sup>1</sup> and for autonomous machines to become an integral part of our daily lives, we still face many technical and nontechnical challenges. In this article, we categorize these challenges in ten areas. For each area, we introduce the current status and roadblocks as well as potential research directions. It is hoped that this will help the

community define clear, effective, and more formal development goalposts for the future.

## AREA 1: THE ON-MACHINE COMPUTE SYSTEM

As opposed to other computing workloads, autonomous machines have a very deep processing pipeline, or computation graph, with strong dependencies among the different stages and strict deadlines associated with each stage.<sup>2</sup> For instance, Figure 1 presents an overview of the processing pipeline of a level 4 autonomous driving system.

Starting from the left side of the figure, the system consumes raw sensing data from mmWave radars, lidars, cameras, and global navigation satellite system (GNSS) receivers and inertial measurement units (IMUs), where each sensor produces raw data at its own frequency. The cameras capture images at 30 frames per second (fps) and feed the raw data to the 2D perception module, the lidars capture point clouds at 10 fps and feed the raw data



**FIGURE 1.** The processing pipeline of an autonomous vehicle. GNSS: global navigation satellite systems; IMU: inertial measurement unit.

to the 3D perception module as well as the localization module, the GNSS/IMUs generate positional updates at 100 Hz and feed the raw data to the localization module, and the mmWave radars detect obstacles at a rate of 10 fps. All of these raw data are then fed to the perception fusion module.

Next, the results of the 2D and 3D perception modules are fed into the perception fusion module at 30 and 10 Hz, respectively, to create a comprehensive perception list of all detected

an output, the control module must still generate a command before the deadline. This could lead to disastrous results as the autonomous machine would then be essentially driving blindly without timely participation from the perception unit.

To minimize the end-to-end latency, one commercial approach is to build a proprietary on-machine computing system to map sensing and computing tasks to the best compute substrates to achieve optimal performance and

the flexibility of mapping various computation graphs from different types of autonomous machines to the same compute substrate, while meeting real-time performance, cost, and energy constraints. Existing CPUs meet the flexibility requirement but fail to meet the performance and energy constraints, whereas other compute substrates, especially accelerators, typically target to meet the performance and energy constraints of one module, for example, the perception module, without optimizing the end-to-end system.

**THIS COULD LEAD TO DISASTROUS RESULTS AS THE AUTONOMOUS MACHINE WOULD THEN BE ESSENTIALLY DRIVING BLINDLY WITHOUT TIMELY PARTICIPATION FROM THE PERCEPTION UNIT.**

objects. The perception list is then sent to the tracking module at 10 Hz to create a tracking list of all detected objects. The tracking list then is fed into the prediction module at 10 Hz to create a prediction list of all objects. After that, both the prediction results and the localization results are received by the planning module at 10 Hz to generate a navigation plan, which then goes into the control module at 10 Hz to generate control commands. These, in turn, are sent to the autonomous machine for execution at 100 Hz.

Hence, every 10 ms, the autonomous machine needs to generate a control command to maneuver the autonomous machine. If any upstream module, such as the perception module, misses the deadline to generate

energy consumption.<sup>3</sup> However, this design evolves through trial and error, and the whole design process takes much time with multiple iterations. This is the same approach that many autonomous machine companies take: they deploy ad hoc solutions to ensure on-time autonomous machine product release. These methods are product specific and hard to generalize to other autonomous machine designs, hence leading to high reengineering costs for each product.

### Challenge

We conclude that the key technical challenge of designing an autonomous machine compute system is to develop an appropriate computer architecture, along with a software stack that allows

### AREA 2: THE SENSING SYSTEM

Autonomous machines gather information about their surroundings through their sensors; sensing is thus the very first stage of autonomous machines' deep processing pipeline. The ultimate objective of the sensing system is to provide accurate, reliable, and comprehensive information about the physical environment for later processing stages to precisely understand the physical world. These three goals (accuracy, reliability, and comprehensiveness) can be achieved through multimodal sensor fusion.<sup>4</sup> However, multimodal sensor fusion presents significant challenges because all sensor data have to be synchronized temporally and spatially all of the time during the operation of the machine.

As an instance of temporal synchronization, take two sensors, a camera and a lidar; their measurements must be made simultaneously so that the machine can fuse the measurements and reconstruct an accurate and comprehensive view of the environment. Without proper time synchronization, the data from multiple sensors could yield an inaccurate, ambiguous view of the environment, leading to potentially

catastrophic outcomes.<sup>5</sup> For an example of spatial synchronization, assume multiple cameras with overlapping field of views are installed at different locations of the autonomous machine. To guarantee that the visual sensing results from these multiple cameras accurately match, the results from each camera need to go through geometric transformations, or calibration, so that they can be projected to a commonly agreed-upon reference point. Currently, on most autonomous machines that fuse multiple sensor modalities, the sensors must be manually calibrated,<sup>6</sup> a slow, labor-intensive, and often error-prone process. Once this calibration has been completed, it is often assumed that it remains unchanged while the autonomous machine is operating. In practice, the calibration is gradually degraded due to the autonomous machine's motion, leading to the need of periodic recalibration.

### Challenge

In an ideal scenario, we should be able to simply plug a sensor into an autonomous machine, and it would just work. Unfortunately, we are still far from this ideal, and any change to the sensor configuration of the autonomous machine must often lead to a redesign of the temporal and spatial synchronization subsystem (for example, a complete recalibration of the whole autonomous machine) or even a modification to the perception and localization systems. We conclude that the key technical challenge for the design of sensing systems is to provide a standardized framework for temporal and spatial synchronization for existing as well as new sensors. Only then will we be able to provide sensor plug-and-play capability for autonomous machines.

## AREA 3: THE PERCEPTION SYSTEM

Perception is essential to any autonomous machine applications where sensory data and artificial intelligence techniques are involved. The final objective of perception is to extract spatial and semantic information from the raw sensing data so as to allow the machine to construct a comprehensive understanding of its operating environment. Such understanding includes the types, positions, headings, speeds, and dimensions of all objects in the environment.

There are two categories of perception for autonomous machines: the deep learning-based approach and the geometry-based approach. The deep learning-based approach is mainly used to extract semantic information and is heavily used in applications such as object detection, scene understanding, segmentation, tracking and prediction, and so on. Within an autonomous machine's perception system, multiple deep learning models are running simultaneously, such as networks for 2D perception, networks for 3D perception, and networks for tracking and prediction. Nonetheless, deep learning-based perception is often the performance bottleneck in autonomous machines. It thus becomes a painful tradeoff between perception quality and compute and energy resource utilization.

The geometry-based approach is mainly used to extract positional and dimensional information of the target objects. A most common instance of geometry-based perception application is the real-time stereo vision used for autonomous machine navigation, obstacle avoidance, and scene reconstruction. Stereo vision allows autonomous machines to obtain 3D structure information of the scene. A

stereo vision system typically consists of two cameras to capture images from two points of view. Disparities between the corresponding pixels in two stereo images are detected using stereo-matching algorithms. Depth information can then be calculated from the inverse of this disparity. Autonomous machines must fuse the positional and dimensional information and the semantic information of the target objects to form a comprehensive understanding of their environments.

### Challenge

We conclude that the key technical challenge for perception is in the development of a general framework to generate a reliable and precise understanding of the operating environment in real time. Reliable and precise perception can be achieved by fusing various perception results, such as 2D semantics, 3D semantics, and 3D geometry, an area still currently under active research. On the other hand, to achieve real-time performance, there are software approaches, such as the compression-compiler co-design method, which combines the compression of deep learning models and their compilation to optimize both the size and speed of deep learning models.<sup>7,8</sup> Hardware approaches such as hardware accelerators for perception modules<sup>9,10</sup> also can be taken. Ultimately, more research is required to determine what combination of software and hardware approaches will be most effective in achieving real-time performance for the perception system.

## AREA 4: THE LOCALIZATION SYSTEM

Fundamental to autonomous machines is localization, that is, ego-motion estimation, which calculates the position

and orientation of an agent in each frame of reference. Formally, localization generates the six-degrees-of-freedom (DOFs) pose, including the three DOFs for the translational pose, which specify the  $\langle x, y, z \rangle$  position, and the three DOFs for the rotational pose, which specify the orientation about three perpendicular axes: yaw, roll, and pitch. Knowing the translational pose fundamentally enables an autonomous machine to plan its path and to navigate, while the rotational pose further lets it stabilize itself.

Localization is highly sensitive to the operating environment, as different environments require different sensors and algorithms. For instance, in outdoor environments, which usually provide stable GNSS signals, using the compute-light visual-inertial odometry algorithm or lidar odometry coupled with GNSS signals achieves the best accuracy and performance. By contrast, in unknown, unmapped indoor environments, a lidar or visual simultaneous localization and mapping (SLAM) algorithm delivers the best accuracy.<sup>11</sup> Nonetheless, different localization algorithms often incur different latencies and, even worse, latency variations. For instance, for visual odometry or visual SLAM, the processing latency often depends on the number of feature points extracted from the current image, which means large latency variations can occur that might impact predictability and safety.

### Challenge

We conclude that the key technical challenge in localization is to develop a standard framework capable of adapting to different operating scenarios by unifying core primitives in various localization algorithms and seamlessly switching among different algorithms

as the autonomous machine navigates through different environments. In addition, this standard framework should provide a desirable software baseline for acceleration, which will minimize processing latency as well as latency variations.

### AREA 5: THE PLANNING AND CONTROL SYSTEM

The planning and control system dictates how an autonomous machine should maneuver. Traditional planning and control systems must include behavioral decisions, motion planning, and feedback control kernels.<sup>1</sup> Specifically, motion planning entails three steps, namely, roadmap construction, collision detection, and graph search. As autonomous machines work with configurations with higher DOFs (such as robotic arms), motion planning will become increasingly complex since the search space will have exponentially increased.

While commonly deployed in commercial autonomous machines, traditional planning and control methods often utilize human-in-the-loop rule-based approaches, where engineers fine-tune the planning and control kernels with available test data. This approach is not only slow and costly, but also it is not robust as all “corner” cases need to be added manually. As more rules must be included, the planning and control system becomes increasingly large and unwieldy and often fails to meet real-time requirements. In addition, rule-based methods suffer from a notorious difficulty to handle the multiagent problem, or the fact that the actions taken by an agent can affect the behavior of other machines in the same environment, hence failing to handle complex traffic scenarios.

Recently, deep reinforcement learning (DRL) for planning and control has been actively researched in many places worldwide. Compared to traditional planning and control methods, inference with DRL incurs low computational requirements during operation, especially for high-DOF configurations.<sup>12</sup> In addition, DRL methods are capable of handling the multiagent problem, allowing autonomous machines to handle complex traffic scenarios. However, model training is the bottleneck for DRL-based planning and control systems, as it requires a vast number of trials to gain enough experience. This is especially the case for complex scenarios, where model training can easily reach millions of steps, with each setup of hyperparameters or reward hypothesis taking hours or even days.

### Challenge

While DRL-based planning and control methods are highly promising, the key technical challenge for planning and control remains the development of a cloud infrastructure that provides sufficient compute power and generates high-quality data for DRL model training. First, to ensure we generate enough high-quality data to train the DRL networks, we need to develop simulation engines that are capable of closely simulating various physical scenarios. Second, to greatly improve algorithm development efficiency, especially for complex scenarios, we need to develop a model training infrastructure that can reduce the training time by orders of magnitudes.

### AREA 6: THE COMMUNICATION SYSTEM

As most autonomous machines are also connected machines, the communication among autonomous machines

as well as that between autonomous machines and the infrastructure is critical to overall performance, safety, and reliability. For instance, cellular vehicle-to-everything (C-V2X) is designed as a unified connectivity platform that provides low-latency communications. It consists of two modes of communications. The first mode uses direct communication links between vehicles, infrastructure, and pedestrians. The second mode leverages commercial LTE or 5G cellular networks to enable vehicles to receive information from the Internet. With the proliferation of 5G technologies, we expect that, in the near future, the C-V2X technology will be extended to all autonomous machines.<sup>13</sup>

Disappointingly, the field experimental results demonstrate that the real-world wireless networks are unstable and unreliable, which can significantly affect the arrival timing of the infrastructure-side data.<sup>15</sup> Further, complex and ever-changing traffic conditions can exacerbate this situation. Specifically, while for direct-communication C-V2X networks, communication latency is less of a problem, the communication bandwidth remains insufficient to sustain the data needed for cooperative autonomous machines. With commercial LTE or 5G networks, while bandwidth is less of a problem, latency variations often lead to missed deadlines where the miss ratio can be as high as 30%, thereby preventing reliable cooperative autonomous driving scenarios.

### Challenge

The key technical challenge of designing autonomous machine communication systems is to develop a highly reliable autonomous machine communication network with bandwidth guarantees and minimum latency variations. There

are some initial directions of research that may achieve these goals, but more research is required. First, we can adjust the priority in the commercial communication networks to provide the highest priority to autonomous machines connections as they are safety critical. Second, a virtual communication backbone that builds on top of both LTE and 5G commercial networks as well as the direct communication C-V2X networks can improve the reliability of autonomous machine communication through intelligent scheduling. Third, we can develop intelligent fusion engines on autonomous machines to dynamically decide whether it is reliable and safe to fuse infrastructure-side data.

### AREA 7: THE CLOUD SYSTEM

Autonomous machines, especially the ones with mobility (such as vehicles, delivery robots, service robots, and drones), require cloud computing support. Such an autonomous machine cloud system must be capable of supporting offline computing tasks, such as deep learning model training or high-definition (HD) map generation, or even online computing tasks such as aggregating all traffic information and broadcasting the information to all traffic participants, especially in the context of cooperative autonomous driving.

Autonomous machine clouds today are mostly application specific, although they share many common requirements. To support different cloud applications, we need an infrastructure to provide basic services, including distributed computing and storage as well as hardware acceleration through heterogeneous computing. If we were to tailor the infrastructure for each application, we would have to maintain multiple infrastructures, potentially leading to low resource utilization, low

performance, and high management overhead. One effective way to solve this problem is to develop a unified autonomous machine cloud infrastructure with consolidated distributed computing and storage services.<sup>14</sup>

Simulation is the cornerstone application in autonomous machine cloud systems, in which autonomous machine functions are exercised in a virtual environment. The architecture of an autonomous machine cloud simulation system is illustrated in Figure 2; the simulator models the sensing mechanism and generates sensory data by simulating interactions between the sensors and environments. Exercised by the simulated sensory data, the autonomous machine software generates control commands to maneuver. We believe that autonomous machine clouds should be simulation centric as simulation consumes outputs from most other cloud-based applications, such as HD map generation and deep learning model training. Also, simulation consumes more than 80% of computing resources in our autonomous machine cloud, making it a low-hanging fruit in terms of a target for optimization. In addition, by moving some of the physical testing to cloud-based simulation, we managed to reduce testing costs by 28 times and improve the testing efficiency by 12 times.

### Challenge

Moving forward, as autonomous machine clouds become more prevalent, two technical challenges have emerged. First, we need to continuously develop and refine distributed computing and storage services to adapt to the needs of various autonomous machine applications, especially to provide the infrastructure support to satisfy the ever-increasing compute and storage

demands for autonomous machines. Second, we need to deploy and optimize basic simulation services on the cloud infrastructure to serve various autonomous machine simulation needs. For instance, autonomous drones and service robots demand different simulation services.

### AREA 8: THE COOPERATION AMONG AUTONOMOUS MACHINES

While traditional autonomous machines utilize only on-machine intelligence, cooperative autonomous machines depend on the cooperation between autonomous machines in addition to the infrastructure. Take autonomous driving for example; the infrastructure–vehicle cooperative autonomous driving approach relies on the cooperation between intelligent roads and intelligent vehicles. This method is not only

safer but also more economical compared to the traditional on-vehicle-only autonomous driving. Based on the progress toward commercial deployment, a three-stage development roadmap has been proposed as follows:<sup>15</sup>

- › stage 1: infrastructure-augmented autonomous driving, in which autonomous vehicles fuse vehicle-side and infrastructure-side perception outputs to improve the safety of autonomous driving
- › stage 2: infrastructure-guided autonomous driving, in which autonomous vehicles can offload all of the proactive perception tasks to the infrastructure in order to reduce per-vehicle deployment costs
- › stage 3: infrastructure-planned autonomous driving, in which

the infrastructure takes care of both perception and planning, thus achieving maximum traffic efficiency and cost efficiency.

To complete these tasks, Figure 3 presents an overview of the cooperative autonomous driving system architecture. It consists of systems on vehicle (SoVs), systems on road (SoRs), the intelligent transportation cloud system (ITCS), and the control center. The SoRs provide local perception results to the SoVs for blind-spot elimination and extended perception to improve safety. Meanwhile, the SoRs process incoming sensor data and send the extracted semantic data to the ITCS for further processing. The ITCS fuses all incoming semantic data to generate global perception and planning information before the control center can dispatch real-time global traffic information,

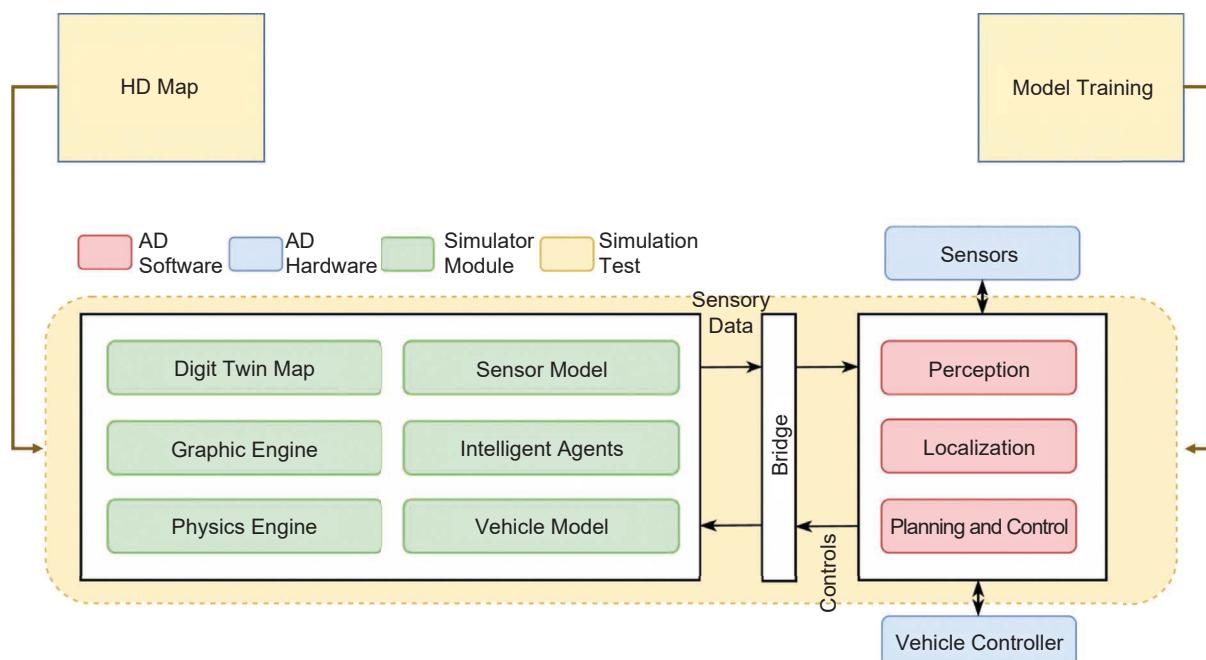


FIGURE 2. A cloud system for autonomous machines. AD: autonomous driving.

navigation plans, and even vehicle control commands to the SoVs to achieve optimal traffic efficiency. This design does not only apply to autonomous driving; it can be generalized to all autonomous machines.

### Challenge

The key technical challenge for cooperative autonomous machines is integration as its full realization relies on all of the technical areas described in the previous sections. To guide the progress of cooperative autonomous machines, the community needs to clearly define the technical specifications and standards of each technical area for each stage of deployment to ensure effective technical integration.

### AREA 9: THE EDUCATION OF ENGINEERS IN THE AGE OF AUTONOMOUS MACHINES

Talent supply is the single most important input for the ubiquitous deployment of autonomous machines. Unfortunately, in the past few years, industry has observed a huge supply–demand gap for autonomous machine engineers.

On the supply side, many students actually have a high degree of interest in autonomous machine technologies but find that their home universities do not provide an adequate education to prepare them to enter this field. On the demand side, it is extremely difficult for autonomous machine companies to hire qualified engineers to fill the open positions, thus forcing many companies to develop internal training programs to prepare incoming engineers. However, since many companies in this field are resource-constrained startups, the need to provide training for incoming engineers imposes an additional burden on their financial situation, leading to an

inefficient allocation of precious funding resources.

To bridge this supply–demand gap, we, as a society, should urgently reshape our engineering education programs and create a cross-disciplinary program to impart students with a technical background in computer science, computer engineering, electrical engineering, and mechanical engineering.<sup>16</sup> On top of this cross-disciplinary technical foundation, a capstone project that will provide students with hands-on experience in working with a real autonomous machine would consolidate the technical foundation.<sup>17</sup> An effective way to perform the capstone project is through cooperative education (coop), a structured method combining classroom-based education with practical work experience, which would require close collaboration between universities and industry.

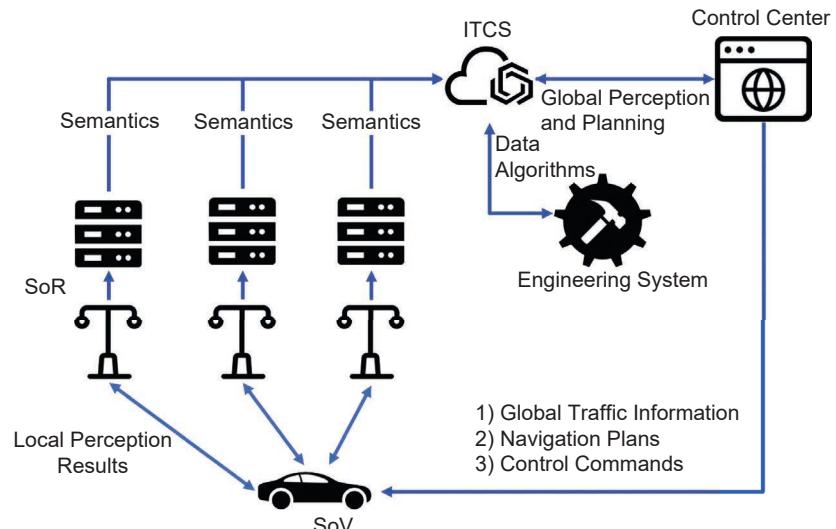
### Challenge

This proposed engineering education reform roadmap is ambitious, and we

have a long way to go to achieve it as it is extremely challenging to establish long-term and stable collaborations among multiple parties. To promote collaboration between industry and academia in emerging fields like autonomous machines, we advocate for governments to provide incentives through grant programs. It is notable that the U.S. government has been playing an active role in driving the U.S. innovation rate.<sup>18</sup> These education and training grants should advance the domestic workforce skills in autonomous machine technologies and consequently improve the domestic economy.

### AREA 10: THE BROADER SOCIETAL IMPACT OF AUTONOMOUS MACHINES

Finally, autonomous machines are expected to completely revolutionize our economy by greatly improving the efficiency in delivery, transportation, manufacturing, and many other sectors. Take last-mile autonomous delivery, for example; the growing



**FIGURE 3.** An infrastructure–vehicle cooperative autonomous driving system. ITCS: intelligent transportation cloud system; SoR: system on road; SoV: system on vehicle.

labor cost may be prohibitive for service providers. In China, a contracted delivery clerk with an annual salary of US\$20,000 can deliver 110 parcels per day, which means that each delivery order costs nearly US\$0.5. This cost is expected to continue increasing as the demographic dividend has reached its end. Deploying autonomous machines for last-mile delivery has proved to be a promising approach to cut these costs by more than half.<sup>19</sup>

Take autonomous vehicles as another example. Traffic inefficiency such as congestion imposes high costs on our society. For the trucking industry alone, the American Transportation Research Institute estimates that congestion costs the United States US\$74.1 billion annually, of which US\$66.1 billion occurs in urban areas. As one of the ultimate goals of autonomous driving is to completely eliminate traffic inefficiency, utilizing autonomous driving technologies on trucking alone can lead to a significantly more efficient national economy.

However, the benefits of improved efficiency do come at a cost. Converting all vehicles in the world into autonomous vehicles would result in an enormous social cost as autonomous vehicles are still very expensive to build.<sup>20</sup> An alternative but cost-efficient solution is cooperative autonomous machines. As infrastructures become more intelligent, more workloads can be offloaded to the infrastructure side, greatly reducing the hardware and energy costs of autonomous machine deployments. It is estimated that this approach would drop the cost of autonomous machine deployment by more than 50%.<sup>15</sup> Besides the impact on the economy, there will also be increasingly more ethical, privacy, and security challenges as autonomous machines proliferate.

### Challenge

One nontechnical but critical challenge is for the whole industry to thoroughly understand the potential impact of various autonomous machines on our economy as well as the costs of large-scale deployments of various autonomous machines. Only with accurate projections and estimations will we be able to provide comprehensive information to policy makers and technology firms. Consequently, policy makers and technology firms are capable of making their policy and technology investment decisions with the ultimate goal of improving human society with autonomous machine technologies. Once we have a clear understanding of the societal benefits of autonomous machines, our society should develop ethical guidelines, standards, and legislations to guide the healthy development of autonomous machines. Hence, an immediate next step is for economists, policy makers, and autonomous machine engineers to work closely together to define economic impact and policy roadmaps to integrate autonomous machines into our daily lives.

**A**fter more than six decades of information technology development, we believe that autonomous machines will completely revolutionize our daily lives and our economy in the coming decade. The impact of autonomous machines on our society is likely to be much deeper and broader than any other information technology revolution that we have experienced in the past decades. To facilitate the rise of autonomous machines, in this article, we have categorized the technical and non-technical challenges in ten areas, and, for each area, we have reviewed the current status, roadblocks, and potential

research directions. Note that the current categorization is by no means comprehensive, as we expect new research areas will emerge as the field evolves. Rather, this article is meant to serve as an initial step for our community to define clear and effective roadmaps to make the age of autonomous machines a reality in the near future. As a next step, we aim to work with technical, educational, economic, and policy experts from academia, industry, and government within each area to form cross-disciplinary teams to address the challenges outlined in this article. □

### REFERENCES

1. S. Liu, L. Li, J. Tang, S. Wu, and J. L. Gaudiot, "Creating autonomous vehicle systems," *Synthesis Lectures Comput. Sci.*, vol. 6, no. 1, pp. i-186, 2017, doi: 10.2200/S00787ED1V01Y201707CSL009.
2. S. Liu, J. Tang, Z. Zhang, and J. L. Gaudiot, "Computer architectures for autonomous driving," *Computers*, vol. 50, no. 8, pp. 18-25, 2017, doi: 10.1109/MC.2017.3001256.
3. B. Yu, W. Hu, L. Xu, J. Tang, S. Liu, and Y. Zhu, "Building the computing system for autonomous micromobility vehicles: Design constraints and architectural optimizations," in *Proc. 53rd Annu. IEEE/ACM Int. Symp. Microarch. (MICRO)*, pp. 1067-1081, doi: 10.1109/MICRO50266.2020.00089.
4. C. Brommer, R. Jung, J. Steinbrenner, and S. Weiss, "MaRS: A modular and robust sensor-fusion framework," *IEEE Robot Autom. Lett.*, vol. 6, no. 2, pp. 359-366, 2020, doi: 10.1109/LRA.2020.3043195.
5. S. Liu et al., "Brief industry presentation: The matter of time—A general and efficient system for precise sensor synchronization in robotic

## ABOUT THE AUTHORS

**SHAOSHAN LIU** is the founder and CEO of PerceptIn Inc., Fremont, California, 94539, USA. His research interests include autonomous driving technologies and robotics. Liu received a Ph.D. in computer engineering from the University of California, Irvine. He is a Senior Member of IEEE, a Distinguished Speaker of the IEEE Computer Society, a Distinguished Speaker of the Association for Computing Machinery, and a founder of the IEEE Special Technical Community on Autonomous Driving Technologies. Contact him at shaoshan.liu@perceptin.io.

**JEAN-LUC GAUDIOT** is a distinguished professor in the Department of Electrical Engineering and Computer Science at the University of California, Irvine, California, 92697, USA. His research interests include multithreaded architectures, fault-tolerant multiprocessors, and the implementation of reconfigurable architectures. Gaudiot received a Ph.D. in computer science from the University of California, Los Angeles. He was the president of the IEEE Computer Society and a Life Fellow of IEEE and the American Association for the Advancement of Science. Contact him at gaudiot@uci.edu.

computing," in *Proc. IEEE Real-Time Embedded Technol. Appl. Symp. (RTAS)*, 2021, pp. 413–416.

6. S. Liu, *Engineering Autonomous Vehicles and Robots: The DragonFly Modular-Based Approach*. Hoboken, NJ: Wiley, 2020.
7. H. Guan et al., "CoCoPIE: Enabling real-time AI on off-the-shelf mobile devices via compression-compilation co-design," *Commun. ACM*, vol. 64, no. 6, pp. 62–68, 2021, doi: 10.1145/3418297.
8. P. Zhao et al., "Brief industry paper: Towards real-time 3D object detection for autonomous vehicles with pruning search," in *Proc. IEEE Real-Time and Embedded Technol. Appl. Symp. (RTAS)*, 2021, pp. 425–428.
9. T. Gao et al., "IELAS: An ELAS-based energy-efficient accelerator for real-time stereo matching on FPGA platform," in *Proc. IEEE Int. Conf. Artif. Intell. Circuits Syst. (AICAS)*, 2021, pp. 1–4.
10. Z. Wan et al., "A survey of FPGA-based robotic computing," *IEEE Circuits Syst. Mag.*, vol. 21, no. 2, pp. 48–74, 2021, doi: 10.1109/MCAS.2021.3071609.
11. Y. Gan et al., "Eudoxus: Characterizing and accelerating localization in autonomous machines industry track paper," in *Proc. IEEE Int. Symp. High-Performance Comput. Archit. (HPCA)*, Feb. 2021, pp. 827–840, doi: 10.1109/HPCA51647.2021.00074.
12. B. R. Kiran et al., "Deep reinforcement learning for autonomous driving: A survey," *IEEE Trans. Intell. Transp. Syst.*, 2021, doi: 10.1109/TITS.2021.3054625.
13. S. Liu, L. Liu, J. Tang, B. Yu, Y. Wang, and W. Shi, "Edge computing for autonomous driving: Opportunities and challenges," *Proc. IEEE*, vol. 107, no. 8, pp. 1697–1716, 2019, doi: 10.1109/JPROC.2019.2915983.
14. S. Liu, J. Tang, C. Wang, Q. Wang, and J. L. Gaudiot, "A unified cloud platform for autonomous driving," *Computer*, vol. 50, no. 12, pp. 42–49, 2017, doi: 10.1109/MC.2017.4451224.
15. S. Liu, B. Yu, J. Tang, and Q. Zhu, "Invited: Towards fully intelligent transportation through infrastructure-vehicle cooperative autonomous driving: challenges and opportunities," in *Proc. 58th Annu. Des. Autom. Conf. (DAC)*, 2021, pp. 1–4.
16. S. Liu, J. L. Gaudiot, and H. Kasahara, "Engineering education in the age of autonomous machines," *Computer*, vol. 54, no. 4, pp. 66–69, 2021, doi: 10.1109/MC.2021.3057407.
17. J. Tang et al., "Teaching autonomous driving using a modular and integrated approach," in *Proc. IEEE 42nd Annu. Comput. Softw. Appl. Conf. (COMPSAC)*, July 2018, vol. 1, pp. 361–366, doi: 10.1109/COMPSAC.2018.00057.
18. S. Liu, "DARPA: A global innovation differentiator," *IEEE Eng. Manag. Rev.*, vol. 48, no. 2, pp. 65–71, 2020, doi: 10.1109/EMR.2020.2975186.
19. B. Li, S. Liu, J. Tang, J. L. Gaudiot, L. Zhang, and Q. Kong, "Autonomous last-mile delivery vehicles in complex traffic environments," *Computer*, vol. 53, no. 11, pp. 26–35, 2020, doi: 10.1109/MC.2020.2970924.
20. S. Liu and J. L. Gaudiot, "Autonomous vehicles like self-driving technologies should start small, go slow," *IEEE Spectr.*, vol. 57, no. 3, pp. 36–49, 2020, doi: 10.1109/MSPEC.2020.9014458.