Machine-Learning-Enabled Cooperative Perception for Connected Autonomous Vehicles: Challenges and Opportunities

Qing Yang, Song Fu, Honggang Wang, and Hua Fang

Abstract

Connected and autonomous vehicles is a disruptive technology that has the potential to transform the current transportation system by reducing traffic accidents and enhancing driving safety. One major challenge of building such a system is how to realize effective and efficient cooperative perception among vehicles, which enables them to share local (raw or processed) perception data with each other or roadside infrastructures through wireless communications. As machine learning techniques become prevalent in autonomous vehicles, particularly in their perception subsystem, we articulate the possibility to design a machine-learning-enabled cooperative perception system for connected autonomous vehicles. Not only are the research challenges in designing cooperative perception presented, but we also focus on how to reduce communication and data processing latency in order to meet the stringent time requirements posed by autonomous driving applications. The article outlines the research challenges and opportunities in designing cooperative perception for autonomous vehicles, leveraging the recent research outcomes from machine learning, feature map quantification, millimeter-wave communications, and vehicular edge computing.

Introduction

Perception for an autonomous vehicle is defined as the ability of a vehicle collecting information and extracting relevant knowledge from sensor data to develop a contextual understanding of the environment. Together with localization and mapping, path planning, decision making, and vehicle control modules, an autonomous vehicle is able to successfully navigate itself on roads. Cooperative perception, on the other hand, enables vehicles to share local perception data with each other (or infrastructures) through wireless communications. One of the prime reasons for developing cooperative perception is the need to maximize the line of sight and field of view of autonomous vehicles. In addition, it could reduce the uncertainty in local object detection results and increase perception accuracy. With increasing situational awareness, cooperative perception is able to expand vehicles' field of view, resulting in safer driving decisions.

Problem Statement

As one of the most challenging tasks in cooperative perception, object detection has been studied for several years. According to the data reported on the KITTI website [1], however, the current best solution to pedestrian detection on autonomous vehicles only achieves a precision of 78.35–83.06 percent. As such, we focus in this article on cooperative object detection to discuss the research challenges and opportunities in cooperative perception for connected autonomous vehicles (CAVs). Traditional cooperative perception is realized at a high level [2], that is, vehicles combine the object detection results shared from others in pursuit of improving their own object detection precision. While it is easy to realize high-level cooperative perception, there is a fundamental flaw associated with this approach. It cannot avoid the issue of what happens if no car senses enough information to detect a critical object.

It was recently proven that more objects can be detected if vehicles share their raw light detection and ranging (LiDAR) data with each other. Cooperative perception relies on vehicle-to-vehicle (V2V) and vehicle-to-infrastructure (V2I) communications; however, due to the limited spectrum allocated for automotive use, it is prohibited to transmit massive amounts of raw data among autonomous vehicles (or between vehicles and roadside infrastructures). The Federal Communications Commission has reduced the spectrum for vehicular communications from 75 MHz to 30 MHz, which cannot support high-volume data transmission. Moreover, fusing data from other vehicles involves data processing, synchronization, and fusion, which will introduce extra latency into the cooperative perception system. Solutions that increase vehicles' perception precision, through cooperative perception, as well as reducing communication and data processing latency are rare in the literature.

Proposed Solution

A useful insight of the current autonomous vehicle's perception system is that modern object detection techniques commonly adopt a convolutional neural network (CNN) [3] to process sensor data. Within a CNN model, sensor data needs to be processed by multi-

Digital Object Identifier: 10.1109/MNET.011.2000560

Qing Yang (corresponding author) and Song Fu are with the University of North Texas; Honggang Wang and Hua Fang are with UMass Dartmouth.

ple convolutional layers, resulting in feature maps that are used by fully connected layers to detect and classify objects. As a different representation of the raw sensor data, feature maps can be viewed as a substitute for the original data to realize cooperative perception. Although the size of feature maps is relatively smaller, a large number of feature maps are generated by a CNN model. It is challenging, if not impossible, to transmit feature maps in their original formats. Therefore, it is crucial to understand how to effectively transmit feature maps within a vehicular network. One possible solution is to leverage the high-speed millimeter-wave (mmWave) communications to realize fast and reliable feature map transmission to realize cooperative perception.

From four different aspects, we outline the research challenges and opportunities pertaining to cooperative perception for CAVs . We first start with how machine learning will enable innovative cooperative perception on CAVs. A key concept is that feature maps generated by the machine learning models on autonomous vehicles can be used to achieve cooperative perception. Following that, we discuss how to quantify the importance of different feature maps, and how to compress and select feature maps to reduce the amount of transmitted data. To facilitate faster data transmission, we then explore the possibility of applying mmWave communication to transmitting feature maps among vehicles. Because it is possible to offload data and computing tasks from vehicles to roadside edge servers, we further describe how to design an efficient vehicular edge computing system to support faster and more reliable cooperative perception in CAVs. Finally, we conclude our work.

Machine-Learning-Enabled Cooperative Perception

For connected and autonomous vehicles, there is no optimal standard or framework that outlines the right level of information sharing for cooperative autonomous driving. It is prohibitively expensive to transmit raw data among vehicles, causing network congestion, packet drops, and large processing delay. Feature maps generated by the machine learning models on autonomous vehicles can be viewed as an alternative representation of the original sensor data, and thus can be transmitted and fused to effectively realize cooperative perception on CAVs. It was proven that feature maps generated on two vehicles can be combined to realize more accurate cooperative object detection [4].

Feature-Map-Based Cooperative Perception

Taking a closer look at how data flows within a CNN [3], we find that when data is processed by a convolutional layer, a feature map is generated, which will be fed into another convolutional layer until it hits the last one. Feature maps are considered the output of their preceding convolutional layer and the input of the succeeding convolutional layer.

When raw data is processed by a CNN, extraneous information is filtered out by the network, leaving behind only essential information for object detection. Data passing through convolutional layers can be viewed as a process of raw data being processed/compressed into feature maps. The major advantage of transmitting feature maps over raw data is that the amount of data transmitted can be significantly reduced. As feature maps extracted from 3D points cloud also contain location information, it is possible to combine the feature maps from different autonomous vehicles. Here, we make the assumption that cars compatible for fusion (e.g., perception systems) manufactured by the same original equipment manufacturing (OEM supplier) will employ the same CNN model for object detection.

Feature maps generated by the machine learning models on autonomous vehicles can be viewed as an alternative representation of the original sensor data, and thus can be transmitted and fused to effectively realize cooperative perception on CAVs. It was proven that feature maps generated on two vehicles can be combined to realize more accurate cooperative object detection.

Feature Map Alignment

The first problem in achieving feature-map-based cooperative perception is how to synchronize feature maps received from different vehicles. After multiple feature maps are received, they must be synchronized in the spatial and temporal domains [5]. By evaluating the similarity between received features, it is possible to align them in the spatial domain; however, research is lacking on how to synchronize feature maps in the temporal domain. Because received feature maps are likely generated at different time instances, it is possible that newly received feature maps were generated at earlier times. Therefore, it is critical to determine the freshness of received feature maps based on local ones. Moreover, feature maps may overlap with each other, so how to fuse the overlapped feature maps becomes an important problem. One possible solution is to assign different weights to different feature maps based on how much new information they can offer.

Feature Map Compression

Although the size of a feature map is smaller than the original data, there are a relatively large number of feature maps generated; for example, SECOND [6] produces 128 feature maps to realize 3D object detection. It is challenging, if not impossible, to transmit feature maps in their original format. Therefore, it is crucial to explore how to effectively compress feature maps to reduce the amount of data transmitted in the network. Previous work has confirmed that feature maps tend to be sparse [7], implying that it is highly possible to compress feature maps to save network bandwidth.

Feature Map Streaming

As data sharing among vehicles is continuous, the feature maps corresponding to the adjacent frames of data will show a certain degree of correlation. Such dependency can be leveraged to reduce both temporal and spatial redundancy, using motion compensation techniques. As a salient function of an autonomous vehicle, real-

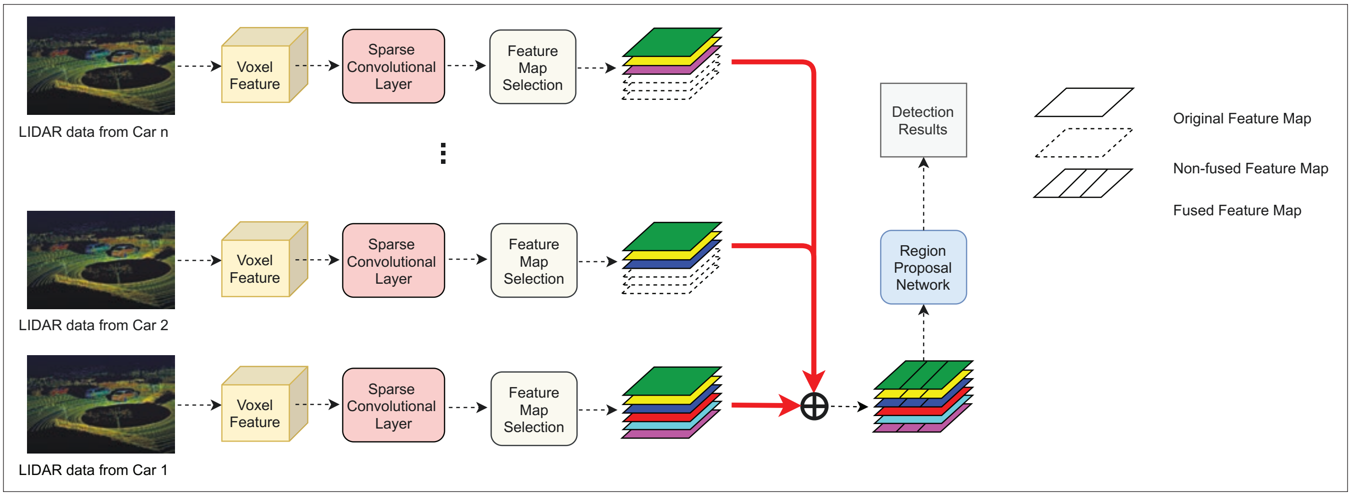

FIGURE 1. System architecture of the feature map selection-based LiDAR data fusion framework.

time object detection results are always available, which could be used to estimate objects' locations in future frames.

Although video compression has been intensively studied in the wireless network domain, the impact of lossy compression on cooperative object detection is not well understood. In particular, the subtle changes of the data in feature maps, due to compression and decompression, may or may not affect the final object detection results. If the information loss occurs in the background regions of a feature map, it would not degrade the final object detection precision. Instead, as the noise/background data is filtered out during the data compression process, the detection results can in fact be improved. On the other hand, if the information loss is related to the objects, worse object detection performance is expected. Therefore, the research challenge lies in how to preserve features pertaining to objects while suppressing those related to the background.

Feature Map Selection in Cooperative Perception

Different feature maps have different impacts on object detection performance; therefore, it is imperative to quantify the importance of feature maps, and then identify and transmit only the most prominent ones among vehicles. As such, it is critical to design a feature map selection-based cooperative perception framework, as illustrated in Fig. 1, where fewer feature maps are transmitted among vehicles to reduce the amount of shared data. As only a subset of feature maps is selected for transmission, the networking delay could be significantly reduced.

Important Feature Map Selection

Previous works have shown that the massive amount of feature maps output by a CNN model's convolutional layers contain lots of zeros [8], implying that not all feature maps are useful. Existing solutions to feature map selection generally seek to structure pruning, which directly removes structured components (e.g., kernels, filters, or even convolutional layers) to simultaneously reduce computation complexity and memory overhead. This is different from the proposed framework, in which the CNN model on vehicles remains the same, as only a subset of feature maps is exchanged among vehicles. Here, we assume that the CNN model on vehicles has already been pruned and is capable of detecting objects precisely. It is therefore important to rank feature maps based on their importance. Only the most important feature maps are selected and shared among vehicles.

The choice of utilizing feature maps for transmission rests on several factors, including data scarcity, wireless channel condition, interference level, distance, and link duration. Although similar solutions can be created to aggregate intermediate feature maps on vehicles, it may take a longer time to process the fused feature maps as they need go through all other filters in the rest of the CNN network. To meet the stringent time requirements posed by cooperative perception in CAVs, it is more efficient to fuse the last-layer feature maps. It is efficient if fewer important feature maps that capture/contain the most prominent information are considered in the object detection task.

Quantifying a Feature Map's Importance

The importance of a feature map can be analyzed from two perspectives: the convolutional kernel that produces the feature map and the value of the information contained in the feature map. As a feature map is usually produced by a 2D kernel in a CNN model, it is intuitive to study the structure characteristics of the corresponding kernel to understand the importance of the produced feature map.

There is a flurry of research on measuring filters in CNNs, including l_1-norm, l_1-regularization, average percentage of zeros (APoZ), group sparsity regularization, and kernel sparsity and entropy [9]. All the aforementioned solutions are designed for network structure pruning, so it is not clear if they are suitable for feature map selection in cooperative perception. As existing feature quantification criteria is very diverse, the question raised here is "what is needed in a good feature criterion in order to obtain the subset of the most relevant features for the problem at hand?". Aiming to realize precise cooperative perception, it is necessary to perform a comprehensive comparison of existing approaches, and design a method considering both the pros and cons of previous ones.

Understanding kernels is the first step in quantifying feature map importance; it is still of paramount importance to explicitly measure the information contained in feature maps. Because feature maps are the most direct reflection of the original data, they must be considered in feature map selection. Entropy plays a central role in information theory as it is proportional to the amount of information to be measured. Therefore, it is possible to use entropy to measure the information contained in a feature map.

Reduction on Data Transmission

According to the design of the proposed feature-map-based cooperative perception, we carry out a preliminary study on how much data reduction can be achieved through feature map compression and channel selection. Figure 2 shows the different amounts of data transmitted between two vehicles when different data processing mechanisms are adopted. As raw LiDAR data are exchanged in Cooper [10], it generates the highest level of data transmission. If feature maps are transmitted, as in F-Cooper [4], the amount of data can be significantly reduced. Furthermore, if compression and/or channel selection are applied, the amount of network traffic can be reduced to around 0.1 MB. The overall data processing time, excluding data transmission delay, for each case is 0.292, 0.281, 0.235, 0.249, and 0.266 s, respectively. The experiment is carried out on a desktop with an Intel I7 7700 CPU, 16 GB memory, a Nvidia Geforce 1060 6 Gb GPU, and Ubuntu 18.04.

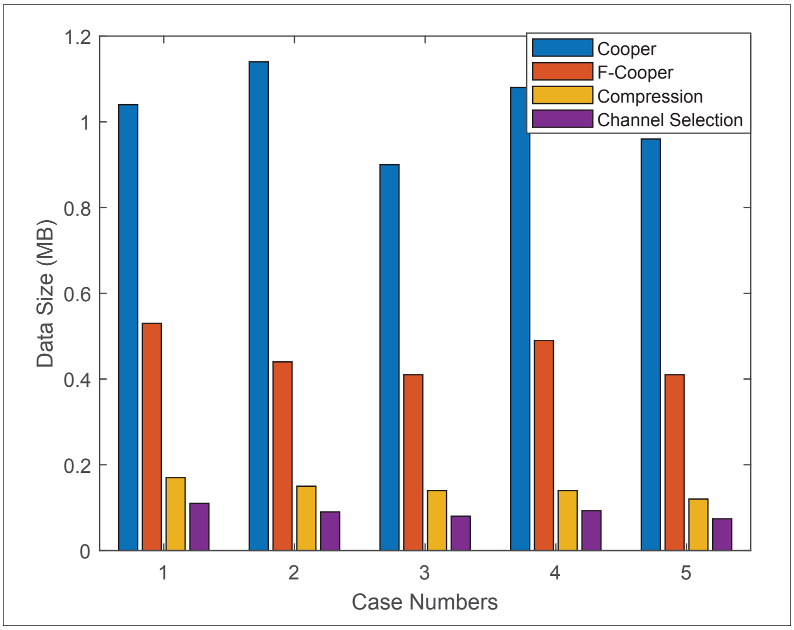

FIGURE 2. Comparison of data volume using different data processing solutions.

Millimeter-Wave Communication for Cooperative Perception

As mmWave communication [11], operating between 10 GHz and 300 GHz, has very large bandwidth, it can achieve multi-gigabit-per-second wireless communications for bandwidth-intensive applications. Therefore, it becomes a perfect solution to massive amounts of data sharing among vehicles to realize cooperative perception. On the other hand, autonomous vehicles can collaboratively achieve accurate perception of surrounding environments, which provides useful information (e.g., receiver's antenna location and communication environments) for designing efficient mmWave communications.

mmWave Vehicular Network

There are some fundamental challenges in terms of mmWave radio propagation, such as higher propagation loss and diffraction. As an mmWave link attenuates over distance, it is better to employ one-hop V2V communications. One-hop mmWave communications allow vehicles to communicate with each other directly, or between a vehicle and a base station (roadside unit), which effectively increases network throughput and improves spectrum efficiency.

It is a challenging problem to deal with multiple V2V pairs in the networks to further improve network throughput and spectrum efficiency for mmWave communications. This problem becomes even more challenging when the high mobility of autonomous vehicles is considered. Due to the large relative velocity, the link duration of a pair of vehicles may not last for a long period. In addition, the constantly changing location of a receiving vehicle may cause a transmitting vehicle to frequently change its beamforming setting to yield reliable communication. Fortunately, because vehicles move on pre-defined roads (available from digital maps) and the roadside infrastructures are stationary, it offers opportunities to design efficient mmWave-based V2V and V2I communication. Especially by predicting a vehicle's future locations based on its instantaneous velocity and the current traffic situation, it is possible to design a continuous beamforming strategy to cope with the mobility of vehicles.

Several studies have investigated the theoretical performance of mmWave communication for autonomous vehicles. However, the feasibility of using mmWave to transmit a massive amount of real-time sensor data among connected and autonomous vehicles has not been thoroughly studied. A research question one might ask is how to design an effective and efficient mmWave vehicular communication system for cooperative perception on CAVs, based on the understanding of the characteristics of mmWave channels.

Sensor-Assisted mmWave Communication

A major difference between autonomous vehicles and conventional vehicles lies in the selection of sensors. To achieve the self-driving function, an autonomous vehicle is typically equipped with various types of sensors, for example, LiDAR, radar, vision camera, thermal camera, ultra-sound, Global Positioning System (GPS), and inertial measurement unit (IMU). To ensure better sensor coverage, multiple sensors of the same type are installed in an autonomous vehicle. These sensors provide real-time sensing data to the processing unit to guarantee that enough information is col-

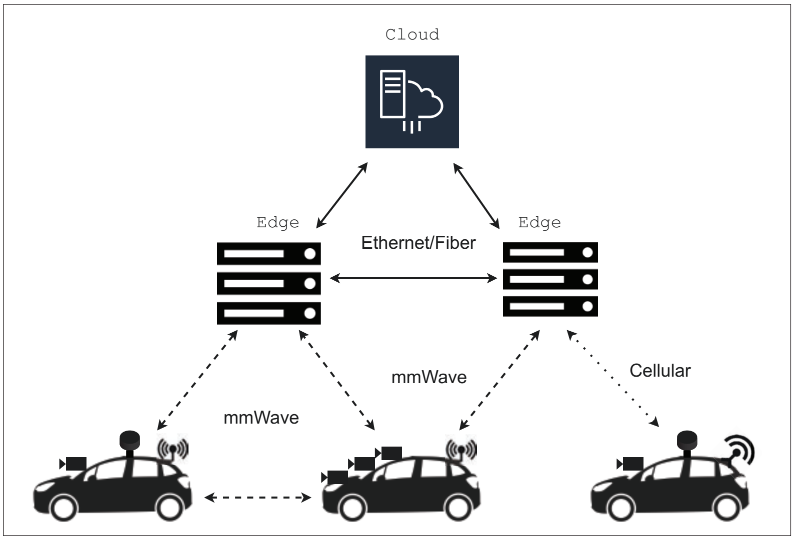

FIGURE 3. Architecture of a vehicular edge system.

lected to allow an autonomous vehicle to successfully accomplish its perception task. Currently, the collected information is barely used for supporting V2V communications; instead, it is mainly used for perception and is discarded/saved after being processed.

The rich semantic information contained in the collected sensor data plays a critical role on setting up a vehicular network. For example, one challenging problem in mmWave communication is beamforming which essentially shapes the radiation pattern of a directional antenna to make sure mmWave signals more powerful and targeted. The state-of-the-art solution to beamforming needs to first transmit a pilot signal over all possible directions, and then measure the corresponding channel qualities to decide the best direction for communication.

Although some works investigate whether GPS information can be leveraged to facilitate effective beamforming for mmWave communications [12], an unexplored problem is to design an effective beamforming solution considering the rich semantic information included in the sensor data. For instance, an advanced image processing model is able to accurately detect the locations of potential receiving antennas from the images captured by autonomous vehicles. Such information could be used to assist the communication system to at least search for a minimum number of directions. Together with other information (e.g., a vehicle's GPS location), a sender vehicle might be able to lock down the target receiver vehicle instantly. Swift beamforming is critical for mmWave communication between vehicles as cooperative perception on autonomous vehicles cannot tolerate large networking delay.

Cooperative Perception in the Vehicular Edge System

In addition to V2V communications, it is important to discuss how V2I communication would affect the performance of cooperative perception on connected and autonomous vehicles. To facilitate effective cooperative perception among vehicles and reduce the onboard computation overhead on vehicles, roadside edge servers bring computation and storage closer to where data is generated, thus reducing response times and saving network bandwidth [13].

Vehicular Edge Computing

The concept behind CAV is as powerful as it is complex, that is, an overlapping group of core technologies are the fundamental building blocks of CAV, including wireless networking, artificial intelligence, computer vision, control, and distributed computing. To tackle this issue, a layered system architecture would allow us to divide a CAV system into different subsystems, each of which composes specific and well-defined parts of the complex CAV system.

We choose a three-layer CAV architecture, as shown in Fig. 3, which contains the perception layer, network layer, and application layer. The perception layer consists of autonomous (and/or conventional) vehicles and roadside infrastructures, equipped with various types of sensors for perceiving surrounding environments. The network layer resides in the middle, and is responsible for data transmission and data processing, also referred to as edge computing. Edge servers are usually deployed on roadside infrastructures (e.g., traffic cabinets) and base stations (e.g., operated by telecommunication companies), or operated by third-party edge service providers. Edge servers are limited in computing and storage resources; therefore, cloud servers are needed to support other key applications, such as simulations, high definition map production, and deep learning model training [14].

Autonomous vehicles can directly communicate with each other if they are using the same wireless communication technology (e.g., mmWave or 5G) [15]. Otherwise, they can coordinate with the aid of edge/cloud servers via relaying traffic or integrating heterogeneous data produced by different vehicles. The application layer is the highes layer, which delivers application-specific services to end users. Edge servers and cloud servers are linked by high-speed networks, such Ethernet or fiber networks, to support delay-sensitive applications.

Computation Offloading

As a killer application of edge computing, CAVs can offload computation-heavy tasks to edge servers to save the precise local computing resources. Most importantly, several computation tasks such as cooperative perception can be better implemented on the edge compared to implementation on individual vehicles. Because the goal of cooperative perception is to let all vehicles exchanging data with each other obtain a more accurate and comprehensive perception of their surroundings, it is straightforward to carry out the data fusion and processing tasks on edge servers and then transmit the processed results to participating vehicles. It would be prohibitively expensive to let each vehicle conduct the same task of data fusion and processing to achieve cooperative perception individually. Thus, it is critical to design an optimized data offloading scheme to facilitate data transmission and data processing in a vehicular edge system.

Resource Task Scheduling

Vehicles tend to form clusters both in town (e.g., between traffic lights) and on highways. Thus, an edge server receives data from multiple vehicles. The number of vehicles that an edge server services can vary dramatically, for example, in rush hours vs. off hours, and in downtown vs. remote areas. One challenge that this class of edge computing faces is the high mobility of vehicles. As a result, the geo-distribution of vehicles is temporary, as is the multi-vehicle cooperation and data processing performed on edge servers. Moreover, edge servers often contain more than one heterogeneous computing unit, with each behaving significantly different in performance on different types of computation tasks. This phenomenon has not been considered in either Linux or third-party frameworks, leading to mismatched resource-task scheduling.

Another research challenge is that edge servers must coordinate the data collection and processing task with each other. For example, it is meaningless to collect data from vehicles whose sensor data are already provided by other vehicles. It is paramount to design an intelligent scheduling algorithm to solicitate feature maps from different vehicles. The design guideline here is to obtain data from the most "unknown" regions that are barely covered. Another opportunity is to integrate the sensor data provided by sensors on roadside infrastructures, for example, cameras or LiDAR sensors on traffic lights, or sensors used in existing road weather information systems.

Conclusions

As an important type of processed data generated by CNN models, feature maps contain sufficient information for autonomous vehicles to accurately detect and classify objects. To realize cooperative perception, sharing feature map data among vehicles is more advantageous compared to sharing the raw sensor data, because the former offers better privacy protection, as well as flexibility on the amount of data to be transmitted. To design a feature-map-based cooperative perception system on CAVs, however, several technical challenges need to be overcome to ensure that the resulting system is reliable and practical. Although some research challenges are identified in this article, including feature map compression, feature map selection, mmWave communications, and vehicular edge computing, they merely pave the way for more advanced solutions to cooperative perception for CAVs. The goal of this article is to open the door for developing new machine-learning-enabled approaches for cooperative perception, and building more efficient and safer autonomous driving vehicles.

Acknowledgments

The work is supported by National Science Foundation grants OAC-2017564, ECCS-2010332, ECCS-2010366, and CNS-1744272. We thank the anonymous reviewers for their valuable comments and suggestions for improving this article.

It would be prohibitively expensive to let each vehicle conduct the same task of data fusion and processing to achieve cooperative perception individually. Thus, it is critical to design an optimized data offloading scheme to facilitate data transmission and data processing in a vehicular edge system.

References

[1] KITTI vision benchmark suite; http://www.cvlibs.net/datasets/kitti/, accessed 19 Nov. 2020.
[2] S.-W. Kim *et al.*, "Multivehicle Cooperative Driving Using Cooperative Perception: Design and Experimental Validation," *IEEE Trans. Intelligent Transportation Systems*, vol. 16, no. 2, 2014, pp. 663–80.
[3] A. Krizhevsky, I. Sutskever, and G. E Hinton, "Imagenet Classification with Deep Convolutional Neural Networks," *Advances in Neural Info. Processing Systems*, 2012, pp. 1097–05.
[4] Q. Chen *et al.*, "F-Cooper: Feature Based Cooperative Perception for Autonomous Vehicle Edge Computing System Using 3D Point Clouds," *Proc. 4th ACM/IEEE Symp. Edge Computing*, 2019, pp. 88–100.
[5] H. Qiu *et al.*, "AVR: Augmented Vehicular Reality," *Proc. 16th Annual ACM Int'l. Conf. Mobile Systems, Applications, and Services*, 2018, pp. 81–95.
[6] Y. Yan, Y. Mao, and B. Li, "SECOND: Sparsely Embedded Convolutional Detection," *Sensors*, vol. 18, no. 10, 2018, pp. 33–37.
[7] B. Liu *et al.*, "Sparse Convolutional Neural Networks," *Proc. IEEE Conf. Computer Vision and Pattern Recognition*, 2015, pp. 806–14.
[8] H. Hu *et al.*, "Network Trimming: A Data-Driven Neuron Pruning Approach Towards Efficient Deep Architectures," arXiv preprint arXiv:1607.03250, 2016.
[9] Y. Li *et al.*, "Exploiting Kernel Sparsity and Entropy for Interpretable CNN Compression," *Proc. IEEE Conf. Computer Vision and Pattern Recognition*, 2019, pp. 2800–09.
[10] Q. Chen *et al.*, "Cooperative Perception for Connected Autonomous Vehicles Based on 3D Point Clouds," *Proc. 2019 39th IEEE Int'l. Conf. Distributed Computing Systems*, 2019, pp. 514–24.
[11] T. S. Rappaport *et al.*, *Millimeter Wave Wireless Communications*, Pearson Education, 2015.
[12] V. Va *et al.*, "Beam Design for Beam Switching Based Millimeter Wave Vehicle-Toinfrastructure Communications," *Proc. 2016 IEEE ICC*, 2016, pp. 1–6.
[13] Z. Ning *et al.*, "Intelligent Edge Computing in Internet of Vehicles: A Joint Computation Offloading and Caching Solution," *IEEE Trans. Intelligent Transportation Systems*, 2020.
[14] S. Liu *et al.*, "Edge Computing for Autonomous Driving: Opportunities and Challenges," *Proc. IEEE*, May 2019.
[15] Z. Ning *et al.*, "When Deep Reinforcement Learning Meets 5G-Enabled Vehicular Networks: A Distributed Offloading Framework for Traffic Big Data," *IEEE Trans. Industrial Informatics*, vol. 16, no. 2, 2019, pp. 1352–61.

Biographies

Qing Yang [M'11, SM'17] (Qing.Yang@unt.edu) is an assistant professor in the Department of Computer Science and Engineering at the University of North Texas, Denton. He received his Ph.D degree in computer science from Auburn University in 2011. His research interests include the Internet of Things, connected and autonomous vehicles, and network security and privacy.

Song Fu [M'08, SM'15] (Song.Fu@unt.edu) is an associate professor in the Department of Computer Science and Engineering at the University of North Texas, Denton. He received his Ph.D degree in computer engineering from Wayne State University in 2008. His research interests include parallel and distributed systems, cloud and edge computing, and machine learning.

Honggang Wang [M'05, SM'10] (hwang1@umassd.edu) is a professor at the University of Massachusetts (UMass) Dartmouth. His research interests include wireless health, body area networks, cyber security, mobile multimedia and cloud, wireless networks and cyber-physical systems, and big data in mHealth. He has also been serving as the Editor-in-Chief of the *IEEE Internet of Things Journal*.

Hua Fang [M'08, SM'17] (hfang2@umassd.edu) is an associate professor at UMass Dartmouth. Her research interests include machine and statistical learning, broadly in E-/M-health, and the Internet of Things.