


Review of Learning-Based Longitudinal Motion Planning for Autonomous Vehicles: Research Gaps Between Self-Driving and Traffic Congestion

Transportation Research Record
1–18© National Academy of Sciences:
Transportation Research Board 2021
Article reuse guidelines:sagepub.com/journals-permissions

DOI: 10.1177/03611981211035764

journals.sagepub.com/home/trr

Hao Zhou¹, Jorge Laval¹, Anye Zhou¹ , Yu Wang¹ ,
Wenchao Wu¹, Zhu Qing¹ , and Srinivas Peeta¹ 

Abstract

Self-driving technology companies and the research community are accelerating the pace of use of machine learning longitudinal motion planning (mMP) for autonomous vehicles (AVs). This paper reviews the current state of the art in mMP, with an exclusive focus on its impact on traffic congestion. The paper identifies the availability of congestion scenarios in current datasets, and summarizes the required features for training mMP. For learning methods, the major methods in both imitation learning and non-imitation learning are surveyed. The emerging technologies adopted by some leading AV companies, such as Tesla, Waymo, and Comma.ai, are also highlighted. It is found that: (i) the AV industry has been mostly focusing on the long tail problem related to safety and has overlooked the impact on traffic congestion, (ii) the current public self-driving datasets have not included enough congestion scenarios, and mostly lack the necessary input features/output labels to train mMP, and (iii) although the reinforcement learning approach can integrate congestion mitigation into the learning goal, the major mMP method adopted by industry is still behavior cloning, whose capability to learn a congestion-mitigating mMP remains to be seen. Based on the review, the study identifies the research gaps in current mMP development. Some suggestions for congestion mitigation for future mMP studies are proposed: (i) enrich data collection to facilitate the congestion learning, (ii) incorporate non-imitation learning methods to combine traffic efficiency into a safety-oriented technical route, and (iii) integrate domain knowledge from the traditional car-following theory to improve the string stability of mMP.

Self-driving cars are around the corner, quite literally. And yet, despite numerous studies (1–7) on the potential impacts of autonomous vehicles (AVs) and connected and autonomous vehicles (CAVs) on traffic flow, a reliable car-following (CF) model describing the longitudinal dynamics of AVs is still lacking. This makes evaluating the impact of AVs on traffic flow challenging. Recent empirical experiments reveal that the existing longitudinal control systems on level-2 AVs are string unstable (8–10), which indicates that small perturbations (e.g., speed fluctuations) tend to grow upstream of a platoon, and eventually lead to full stop-and-go motions. Those empirical findings are surprising, and indicate that AVs might cause more traffic congestion even than human drivers. The results also distinguish from the successful design or implementation of the string-stable longitudinal controller in the literature, including both adaptive cruise control (ACC) and cooperative adaptive cruise control (CACC) algorithms (11–15). It is conjectured here that

the gap between the practice and the theory may result from: (i) the longitudinal control of level-2 AVs, also known as ACC, not factoring string stability in its design; (ii) in real-world scenarios, some other issues (e.g., safety, efficiency, comfort, or user acceptability) being weighted more than string stability performance, thus the controller will suppress the string stability properties to satisfy other performance metrics (16); and (iii) the hardware equipment (sensing devices and actuators) not being capable of realizing the string-stable control command. The rough and choppy measurements, and the slow-response actuator installed on economy daily cars, require the control command to be heavily filtered before

¹School of Civil and Environmental Engineering, Georgia Institute of Technology, Atlanta, GA

Corresponding Author:

Jorge Laval, jorge.laval@ce.gatech.edu

Table 1. Latest Advanced Driver-Assistance System Technologies from Major Automakers in 2020

Automaker	Technology	Sensors	Description
Tesla	Autopilot and full self-driving (FSD)	Radar + 8 cameras	Traffic-aware cruise control (18)
Nissan	ProPILOT 2.0	7 cameras, 5 radars and 12 sonars	Incorporates 3D high-definition map (23)
Toyota	Full-speed range dynamic radar cruise control (DRCC)	Radar + camera	Work in full-speed range (24)
Honda	Honda sensing	Radar + camera	Collision mitigation braking, speed signs (25)
General Motors	Adaptive cruise control (ACC) with camera	Camera + radar	ACC is based on camera
Ford	Intelligent ACC	Radar + camera	Automatically adapt to speed limit signs (26)
Audi	ACC with stop-and-go	Two radars + camera	Stop-and-go (27)
BMW	ACC with stop-and-go	Radar + camera	Stop-and-go and speed limit compliance (28)

being exerted on the vehicle (otherwise the vehicle would behave in an undesirably jerky manner), which makes string stability not achievable. Given the undesired string unstable ACC, it is possible that current AV systems might induce more instability than human drivers, which could induce more traffic congestion and emissions. From a traffic perspective, there is a critical need for a deeper understanding of AVs' longitudinal behaviors to predict their impact on traffic congestion.

Meanwhile, the current AV technology is fast evolving thanks to the recent advancements in computer vision and machine learning. Notably, we are witnessing a fundamental shift from the traditional radar-based ACC, which relies solely on radar (17), to camera-included advanced driver-assistance systems (ADAS). The transition is reasonable and as expected, because the traditional radar-based ACC has a limited functionality from its pure reliance on the radar sensor and hard-coded human-crafted rules. Additionally, the inherent structure of radar-based ACC may lead to issues such as: (i) inability to adapt to variable speed limits, respond to the ambient traffic proactively, or predict upcoming incidents, (ii) inability to navigate in stop-and-go traffic because of the limitations in detecting slow-moving or still objects, and (iii) to alleviate the traffic oscillations, the hard-coded CF rules also require more human efforts in examining and tuning the controller.

This shift from radar to cameras can be game-changing because vision opens the gate for incorporating more machine learning methods, such as mMP, for planning. The leading company, Tesla, is famous for its camera-based autonomy solution and its latest full self-driving (FSD) function features "traffic-aware cruise control" (18). Starting from May 2021, Tesla completely abandoned radar on new releases of its FSD software (19). Although FSD's cruise control demonstrates multiple intelligent features, there is no reliable evidence to show whether its longitudinal motion planning is powered by neural networks or the traditional rule-based

ACC with extra augmentations. Recently, many other automakers have been catching up and also starting to integrate cameras into the longitudinal control module. A brief summary could be seen in Table 1. In general those automakers adopt a similar ADAS, which adds camera for lane keeping and collision avoidance, and enables low-speed cruise control in stop-and-go traffic where a single radar often fails. General Motors (GM) and Nissan seem to be slightly different. Instead of using radar, GM's current ACC function is reportedly only using camera (20), and its upcoming Super Cruise (21) would be a hands-free function using LiDAR maps of highways. Nissan (22) has delivered level-3 autonomous driving using a complex suite of sensors similar to that of Tesla. Nissan also claims to be the first automaker that incorporates the three-dimensional high-definition map.

On the other hand, although there exist hundreds of AV automakers, there are far fewer AV service providers. The major ADAS service providers with their major customers and collaborators are summarized in Figure 1. More detailed information on service providers of ADAS and other AV technologies are attached in the Appendix. It indicates that, despite the many different brands of AVs, their impact on traffic flow is likely to be similar to each other.

While the level-2 market AVs are proprietary and no explicit knowledge about their longitudinal control methods is available, the self-driving technology companies/institutions have been more transparent and exhibited a clear goal to achieve and adopt the mMP. Waymo published its feature-engineering mMP approach in Bansal et al. (30). Remarkably, an end-to-end mMP model was recently open-sourced by Comma.ai, an aftermarket self-driving company which retrofits regular cars with a mono-camera phone. Similar self-driving service is also seen at Mobieye (31), part of Intel, which helps regular cars to function as AVs with only a single camera device. Similarly, many other self-driving technology companies




Figure 1. Main suppliers and customers of advanced driver-assistance systems (ADAS).

Source: Wu (29).

have published their datasets which indicate mMP methods toward the longitudinal autonomy. Readers are referred to Scale (32) for a full list of those public self-driving datasets, which are filtered by data type, traffic scenario diversity, and annotation. On the other hand, a plethora of research papers (33–37) have been proposed to accomplish mMP using different learning approaches.

With all that being said, it is highly possible that mMP will be the future of AVs, for both level-2 commercial vehicles and the higher-level FSD cars according to the definition of the Society of Automotive Engineers. As its impacts on traffic congestion are essential and have not yet received enough attention, an in-depth review is necessary to understand the state-of-the-art mMP methods, with the purpose of promoting more traffic-friendly AVs in the long run.

There already exist some review works on AV planning algorithms in the literature, but their focus is not related to traffic congestion or mMP methods. For example, Babak et al. (38) is limited to the engineering perspective only, focusing on sensors and embedded systems for AVs. Tesla (17), Katrakazas et al. (39), and Paden et al. (40), in the robotics literature, discussed the traditional motion planning approaches like graph search, trajectory optimization, and optimal control methods, which are out of the scope of this study. Quite a few reviews focused on the rule-based AV control, especially for CACC (41). Attempts to consolidate more relevant

studies on mMP of AVs are available. Ni et al. (42) introduced the development of AVs and basics of deep learning methods, as well as summarizing recent research on theories and applications of deep learning for AVs. However, they aimed to identify challenges and solutions in learning algorithms and took an overview from the vehicle perspective. A summary or discussion from the system perspective, such as the impact of mMP on traffic congestion, has not been presented. Similar conclusions can be drawn from the reviews by Schwarting et al. (43) and Yurtsever et al. (44). To the best of the authors' knowledge, the only work that overviews learning-based AV control methods from artificial intelligence (AI) in the field of transportation engineering is Di and Shi (45). Nonetheless, that survey was focused primarily on how to deal with interactions between AVs and human-driven vehicles, especially by reference to academic works.

Compared with the existing review papers on AV control, the aim of this study is to provide a comprehensive outlook to consolidate the existing knowledge base of upcoming mMP of AVs and their impacts on traffic congestion. Specifically, this review paper aims to answer the following questions:

- Data: Do existing self-driving datasets contain congested scenarios? Do they include the necessary features/labels to train a congestion-mitigating mMP?
- Learning method: What are the potential strengths and weaknesses of the typical learning methods in their impact on traffic congestion?
- Domain knowledge: How could expert knowledge of traffic flow help the AI community build the congestion-mitigating mMP?

To this end, the paper is organized as follows. The second section introduces available open datasets for AV development; the third section summarizes learning methods for AV control; the fourth section discusses the major limitations and challenges arising from these previous works; the fifth section proposes how to utilize traffic domain knowledge to leverage current mMP, and the final section presents the discussion and outlook based on this review work.

Datasets for mMP

A typical framework in modern autonomous driving systems is shown in Figure 2. Among those pillars, the mMP in this paper falls into the driving policy/path planning module. Following the pipeline, the related components of training data, model input and output are reviewed, as well as the learning methods for mMP.




Figure 2. Fixed modules in modern autonomous driving systems. Source: Talpaert et al. (46).

Available Open Datasets

Two recent studies (47, 48) provided good reviews of the existing open datasets, which covered data scale, contents (camera or LiDAR, object annotation), road scenarios (urban streets or highway), weather conditions and test vehicle type. Most of the current open datasets are designed to assist computer vision development, even leaving out the some information (e.g., acceleration, trajectory data) required to mimic human driving. From the perspective of traffic congestion, Table 2 summarizes the datasets including the position information that is necessary for learning mMP. This paper also shows specific concerns of the driving scenarios and traffic conditions which are certainly related to mMP model.

Among the currently existing datasets, nuScenes (54) and HighD (62) have shown some consideration of congestion. The nuScenes dataset collected data from Boston and Singapore, two cities that are known for their dense traffic and highly challenging driving (242 km traveled at an average of 16 km/h). The HighD dataset was recorded at six different locations near Cologne, Germany. However, the authors are not aware of any studies that have used nuScenes or HighD to train an autonomous driving system. Recently, more AV companies from the industry, like Waymo and Lyft, have released some open datasets. Waymo's dataset (49) does not provide direct information on trajectory, one needs to derive it using kinematics information. Lyft's dataset (65) does not cover congestion scenarios. Tesla has not revealed any plan to publish its dataset yet, but the authors conjecture that with their large deployment of vehicle fleets it would be highly possible to gain enough congestion data. Remarkably, the L3Pilot dataset will record the autonomous driving behavior and the trajectories of 13 OEM autonomous driving systems, which includes 1,000 drivers and 100 cars in various driving conditions (i.e., different weather and traffic conditions) across 10 countries in Europe. The comprehensive coverage and enriched features of the L3Pilot dataset can

significantly enhance the research on autonomous driving. However, the L3Pilot autonomous project is still ongoing and the corresponding dataset is not yet available to the public. Thus, the current overall situation indicates the lack of consideration of congestion in both academia and industry. Next-Generation Simulation (NGSIM), an open dataset consisted of two-dimensional trajectories, has been widely used in CF studies for decades. Different from the datasets from the AV industry, the traffic density in NGSIM often varies significantly and covers both full states from free flow to traffic jams. It also exhibits a high degree of vehicle interaction near traffic bottlenecks like on-ramps or off-ramps. The diversity of driving scenarios and the interaction among vehicles makes NGSIM especially valuable for learning driving behaviors under congestion. However, it does not provide any image or LiDAR data compatible with sensors for AVs. Moreover, the OpenACC dataset (64) provides the highway trajectory data of multiple vehicles driven by different commercial ACC systems. However, similar to the NGSIM dataset, the OpenACC dataset does not provide any video data or vehicle sensor recordings that be can leveraged in end-to-end mMP.

A general issue in those open datasets is that it is unclear whether those miles were driven by human drivers, traditional ACC controllers, or new mMP models. It becomes a major limitation when researchers attempt to reverse-engineer those current mMP models or simply use the data for training. It might also explain why the applications of those open datasets to transportation studies are still very limited. Overall, the current datasets from the self-driving industry are very limited for analyzing the impact of mMP method on traffic congestion. While more and more commercial ACC products are expected to be equipped with mMP in the future, it would be beneficial for research purposes if car companies were to share their driver data.

Simulator Datasets

While it is costly to collect data from the real world, high-fidelity driving simulators have also been developed to train AVs. CARLA (66) and TORCS (67) might be the most popular open-source simulators for autonomous driving research. Related studies based on those simulators can be found in Chen et al. (68), Panwai et al. (69), Codevilla et al. (70), Mirowski et al. (71), Tan et al. (72). CARLA can define diverse sensor suites and is also able to generate congested traffic scenarios. A specific method of transferring driving policies from simulations to the real world was shown in Müller et al. (73). Note that those simulators also make the reinforcement learning (RL) method feasible by providing an interactive environment for agents to learn.

Table 2. Open Datasets and Simulators for Training Autonomous Driving Systems

Dataset	Data Type		Scenario	Traffic	Highlights and comments
	T	C			
Waymo (49)	Y	Y	U	Light	Kinematics derived from speed
ApolloScape (50)	Y	Y	U	Light and dense	Cover different traffic densities
KITTI (51)	Y	Y	H, U	Light and dense	The first autonomous vehicle dataset, mainly used for vision
BDDV (52)	Y	Y	H, U	Light and dense	A large-scale diverse driving video dataset with comprehensive annotations
Udacity (53)	Y	Y	U	Light	Driven by ACC, in 2016
nuScenes (54)	Y	Y	H, U	Dense	Boston and Singapore, including congestion
Ford (55)	Y	Y	H, U	Intermediate	Including car-following in congestion
Argoverse (56)	Y	Y	H, U	Intermediate	Annotation and labels included in the video
NGSIM	Y	N	H, U	Light and dense	Mostly used by car-following model studies
Comma.ai (57)	Y	Y	H	Intermediate	Driven by ACC and human drivers
Brain4Cars (58)	Y	Y	H, U	Unclear	Behavioral label
CityScapes (59)	Y	Y	U	Dense	Diverse real-world driving scenes with high-quality annotation
Oxford RobotCar (60)	Y	Y	U	Light and dense	Diverse traffic conditions for the whole year in Oxford, UK
UAH (61)	Y	Y	H	Intermediate	Driving behavior analysis with IOS app
HighD (62)	Y	N	H	Light and dense	High-resolution drone data with extracted features
L3Pilot (63)	Y	Y	H, U	Light and dense	First comprehensive test of OEM(Original Equipment Manufacturer) self driving systems in European Union
ACC data (9, 10, 64)	Y	N	H	Intermediate	Trajectories of recent ACC car models

Note: H = highway; U = urban; T = trajectory data; C = camera data; Y = yes; N = no; ACC = adaptive cruise control.

Both academia and industry have been using simulator datasets to test AV software and hardware. For example, in academia, developers from CMU and MIT used TROCS and Talos simulators, respectively, to test their algorithms in simulation before porting them to the vehicle for practical road test (74, 75). Recently researchers have used simulated LiDAR data to develop and test algorithms for AV off-road ground navigation using the MSU autonomous vehicle simulator (76). To supply the critical events and corner cases for the evaluations of AVs efficiently and effectively, Feng et al. (77) leveraged RL algorithms to generate naturalistic adversarial critical events in CARLA to test the safety performance of AVs. In industry, simulator datasets have been used by car manufacturers not only to eliminate modeling errors and validate control systems for AVs (78), but also to evaluate the powertrain performance and the analysis of energy consumption of AVs (79, 80). Waymo (79, 81) and Uber (82) developed simulator platforms to generate realistic scenarios from their real-world datasets to improve the safety and performance of AVs.

For the impact on traffic congestion, similar simulation-based methods can be adopted to generate more driving scenarios related to the traffic efficiency, besides the safety-oriented experiments. However, even though the simulation-based method is efficient in generating supplementary data, simulating realistic behavior of human drivers in a complex traffic environment remains a difficult task, because surrogate models used in simulation will inevitably induce model bias and

over-simplified behaviors. The simulation environment constructed with such a surrogate model can lead to undesired and biased performance measure of AVs. Alternatively, we could use a simple CF model known to be string stable to train a string-stable mMP. Despite the potential benefits of simulator datasets, studies incorporating them to develop a congestion-mitigation mMP model have not been reported. Some studies from the transportation research domain might be close (83, 84), using a simple traffic simulator to train a single AV to stabilize mixed traffic. However, the studies using more high-fidelity driving simulator data to investigate string-stable mMP models have not been found.

Learning Method

Behavior Cloning

A simple yet effective learning method for mMP is to map model inputs to outputs directly, which can be represented as a function mapping from the input features s (e.g., video frames, figure annotations, kinematic information of ambient vehicles, etc.) to the output action space a (e.g., vehicle speed, acceleration, and steering angle, etc.): $F(s) \rightarrow a$. This method is referred to as behavior cloning (BC), a subset of imitation learning. The classical framework of BC methods for mMP can be classified into three categories: end-to-end learning, mid-level learning, and mixed (hybrid) learning approach. These methods are specifically discussed below:

(i) End-to-end mMP. The end-to-end learning approach behaves similarly to a black box, which takes in the raw video data and outputs the longitudinal vehicle control command (e.g., speed, acceleration, throttle response). Even though the end-to-end approach preserves the advantages of self-optimizing and requiring less manual effort in implementation, it does confront difficulties and challenges in capturing and processing crucial features from raw video frames. Specifically, the video data in traffic congestion would contain multiple clusters and pose great difficulty to image processing and feature extraction. In addition, the congestion data may contain undesired noise or become excessively random for neural networks to learn, which might trigger underfitting or overfitting issues. The strategies reported in the existing literature solely rely on two categories of neural networks: convolutional neural networks (CNN) and recurrent neural networks (RNN). For instance, Kim and Canny (85), Bojarski et al. (86), Chen and Huang (87), and Sharma et al. (35) utilized deep CNNs concatenating with multiple fully connected layers to predict the vehicle steering wheel angles, which demonstrated a decent performance in the real-world driving scenario. Moreover, researchers are also contributing to the vehicle longitudinal command. Considering the spatial-temporal characteristics and the memory impact of vehicle longitudinal trajectories, the long short-term memory (LSTM) or gated recurrent unit (GRU) augmented deep CNN (52, 88, 89) are applied to artificially forget or remember the historical frame features to improve the accuracy of vehicle longitudinal commands (i.e., speed, acceleration) prediction.

(ii) Mid-level learning. The mid-level learning method is more interpretable compared with end-to-end learning approach because of its explicit hierarchical structure. The first segment of mid-level learning is to extract the useful CF features (e.g., inter-vehicle spacing, relative speed, lane position, etc.) using computer vision algorithms, then the second segment correspondingly retrofits the CF model with specific neural network. Remarkably, Zhou et al. (90) showcased the effectiveness of an RNN-based CF model in capturing the traffic oscillation characteristics, which provides an insight on including RNN (e.g., LSTM, GRU) in the deep neural network to retrofit the CF behavior in congested traffic condition. Moreover, some studies (91–94) have demonstrated that by arranging the kinematic information of multiple neighbor vehicles in Laplacian-like feature matrices or tensors and applying graph convolution network to seize the inter-dependency and social pooling of data, performance in predicting the states of ego vehicles could be improved. This phenomenon indicates that features with higher dimension and organized in connected structure might lead to higher accuracy. Under this circumstance,

it is also significant to evaluate those hand-crafted features with regard to the model accuracy and parsimony, such that a trade-off can be achieved between model complexity and accuracy.

(iii) Mixed (hybrid) learning approach. As including more useful features in the tensor can boost the prediction performance, some studies have also included another sub-task (e.g., semantic segmentation, image augmentation) to extract those useful features in the training process or incorporated other information (e.g., vehicle kinematic states, ambient traffic information) into the end-to-end learning to improve the model accuracy. For instance, George et al. (34), Yang et al. (95), Hsu et al. (96), and Li et al. (97) pooled the vehicle kinematic information with the features obtained from video frames using concatenating layer to enhance the prediction of steering angle and acceleration. Xu et al. (52) conducted a semantic segmentation aside of the longitudinal and lateral end-to-end learning, and added the loss function of semantic segmentation to the driving loss function of end-to-end learning to reinforce the prediction accuracy. It was found that the simultaneous learning of semantic segmentation could outperform both the end-to-end and mid-level learning methods.

Remarkably, the BC method has gained wide popularity within the industry. Waymo's research paper (30) reported that, even with 30 million examples and mid-level input and output for motion planning, a pure BC method is not sufficient to train a safe AV. To tackle this, they synthesized more "corner" cases through adding perturbations to the normal driving data. However, it is conjectured that it might not lead to much difference since the longitudinal motion planning under normal driving scenarios is not strengthened by "corner" cases. Although Tesla has not published any official research documents on its motion planning technology, from its investor conference event in April 2019 (98), one could speculate that Tesla most probably adopts the BC method as well, and the supervised learning model is evolving with the large deployment of vehicle fleets on the roads. Currently, Tesla is adopting a feature-engineering approach rather than an end-to-end method. Evidence can be found from the videos (99) on the Autopilot official website, in which entities such as vehicles, traffic lights, or cones are all labeled and annotated separately. Moreover, at the Scaled Machine Learning conference in February 2020 (100), Tesla revealed the neural network architecture applied in the FSD, from which it appears that Tesla is applying a HydraNet for pooling different neural networks which conduct different tasks of perceptions and predictions (e.g., labeling, annotation, semantic segmentation, per pixel depth prediction) but share the same backbone. Correspondingly, the HydraNet fuses the information from all cameras

and radars to create a bird's-eye view for navigating the vehicle.

IRL and GAIL

Another pipeline of imitation learning is to recover the implicit reward function of human driving using inverse reinforcement learning (IRL). IRL defines the cost function of a trajectory c_θ and maximizes the probability of expert demonstration:

$$p_\theta(\tau) = \frac{1}{Z} \exp(-c_\theta(\tau)) \quad (1)$$

where τ is a state-action trajectory, and Z is the integral of $\exp(-c_\theta(\tau))$ over all trajectories that are consistent with the environment dynamics (101). The parameters θ are optimized to maximize the likelihood of the demonstrations. If the cost function is learned, one can simply use RL to find the policy that behaves identically to the expert. The first IRL study for autonomous driving was that of Abbeel and Ng (102), which proved that it is possible to “guess” the cost function for some simple task like highway driving by approximating it with a linear combination of some hand-selected features. Related works can be found in Sadigh et al. (103), González et al. (104), and Sharifzadeh et al. (105). However, linear assumption of the reward function will lead to ill-posed problems, because the probability of expert behaviors can be maximized by many different parameters θ . Thus IRL was extended to maximum entropy ILR by Ziebart et al. (106). However, IRL methods are typically computationally expensive in their recovery of an expert cost function and generally require RL in an inner loop.

Noticing the immense computational cost in recovery of the true reward policy, Ho and Ermon (107) found that human driving behaviors can be mimicked directly using generative adversarial imitation learning (GAIL) without discovering a cost function first. GAIL trains the self-driving policy π_θ to perform expert-like behaviors by rewarding it for “deceiving” a classifier D_ϕ that discriminates between the policy and expert state-action pairs. Suppose driving is a sequential decision-making task following a stochastic policy $\pi_\theta(s, a)$, which maps an observed road condition s to a distribution over driving actions a . Sample a set of simulated state-action pairs $\chi_\theta = (s_0, a_0), (s_1, a_1) \dots (s_T, a_T)$ using parameterized policy π_θ , and the expert behavior pairs χ_E from π_E , the GAIL objective is:

$$\max_{\phi} \min_{\theta} V(\theta, \phi) = \mathbb{E}_{(s, a) \sim \chi_E} [\log D_\phi(s, a)] + \mathbb{E}_{(s, a) \sim \chi_\theta} [1 - \log D_\phi(s, a)] \quad (2)$$

In a recent work, GAIL was applied to the task of autonomous driving on highway scenario using NGSIM

dataset (108). The result shows that the recurrent GAIL is surprisingly able to capture many desirable properties consistent with real trajectories. Bhattacharyya et al. (109) extended GAIL to multi-agent learning for highly interactive driving cases. Although the methodology of GAIL is sound, there do not appear to be more follow-up studies from the academic community or industry.

Reinforcement Learning (RL)

The success of imitation learning largely depends on the availability and distribution of labeled data, which are costly to collect. To circumvent this problem, another stream in mMP is working on the non-imitation method, RL, which follows a pipeline as shown in Figure 3. Since RL methods need expert-designed reward functions, they can be designed according to the basic driving rules for autonomous driving, such as gaining faster speed and avoiding collisions. Pan et al. (110) used RL to train an autonomous driving policy with a pre-defined reward function encouraging higher speed and penalizing crashes. In more recent work, Chen et al. (111) implemented several deep RL methods and showed good driving performance with dense surrounding traffic. Guo et al. (112) used the RL method to learn the longitudinal motion planning for AVs to reduce fuel consumption as well as to maintain acceptable travel time. Shalev-Shwartz et al. (113) applied multi-agent RL in a highly interactive merging case to generate a set of feasible trajectories and then feed a hand-designed cost function to the trajectory planner to select the most smooth and safe trajectory, which makes the longitudinal motion planning no longer a pure BC process. DeepTraffic (114), a simulation and deep RL environment developed by MIT, has also shown the success of RL in navigating AVs on a congested seven-lane highway. Other similar studies based on RL and traffic simulators can be found in Sallab et al. (115), Kendall et al. (116), and Liang et al. (117).

It is worth noting that, in the RL context, the model input also plays an important role because it directly determines the state space that an RL agent can observe. Chen et al. (111) reduced the state complexity through feature representation based on the raw image, which makes the problem more tractable and computationally efficient. Despite some studies using RL to stabilize mixed traffic in a loop (118) or near the merging areas (83), there does not yet appear to have been any success in learning a string-stable mMP for single AVs.

In summary, the involvement of the major motion planning methods is shown in Figure 4, which depicts the transition from traditional rule-based methods to the state-of-the-art machine learning methods. Note that most learning methods fall into the range of BC, and




Figure 3. Architecture of training autonomous driving in simulation using reinforcement learning (RL).

Source: Chen et al. (111).




Figure 4. Trend of major learning methods for mMP of autonomous vehicles (AV).

although many alternative learning methods for BC have been proposed in the literature, the leading AV companies still stick to BC (30, 98). Here RL is not considered as a BC method since RL does not directly learn from expert demonstration. It does not require large amounts of data but a high-fidelity simulator. Also, the performance of RL heavily depends on the human-designed reward functions that govern the training process and resulting policies.

Limitations of mMP

Based on the previous review, this section will discuss the current limitations of mMP research with regard to its impact on traffic congestion.

Systematic Lack of Training Data

Datasets that can completely cover regular driving scenarios are still unavailable, let alone the “corner” cases

that threaten the robustness of mMP. No driving data were found for multi-lane highways, on-ramp and off-ramp bottlenecks, or generally congested traffic conditions. Since most neural network methods cannot generalize well to unseen situations, the authors believe that the incompleteness of datasets might lead to biased or even unpredictable CF behaviors. Issues of such limitations in biased datasets were also discussed by Codevilla et al. (119).

Incomplete Feature Representation

While perception modules can extract human-interpretable features as model inputs for mMP, those hand-selected features may not fully capture all the influencing factors for driving decisions. For example, the specific location information might be totally ignored in model input. From the industry, no information has been revealed about whether the localization results are incorporated into motion planning. While human drivers respond to different locations with varying driving behaviors, such as the “relaxation” phenomenon discovered by Laval and Leclercq (120), we still do not know whether mMP will react differently in traffic bottlenecks, such as on-ramps or off-ramps.

Codevilla et al. (70) and Sauer et al. (121) conditioned the BC with high-level command input for intersections. The included high-level commands are able to resolve ambiguities in the mapping from single image input to low-level commands (steering and speed). It is argued here that in highway driving, such ambiguities will also arise between the exiting and non-exiting vehicles (120). Thus it would be worthwhile to incorporate driving intention in motion planning for AVs. However, only Tesla (98) has reported a related project to infer the lane

change intention of leading vehicles and integrate it for motion planning.

Limitations in Learning Algorithms

According to Kuefler et al. (108), the BC method has been successfully used to produce driving policies for simple scenarios such as CF on freeways. However, Wheeler et al. (122) and Lefèvre et al. (123) reported different results when applying BC to nuanced states with little or no experience, showing that BC can only produce accurate predictions up to a few seconds. Their results indicate that BC usually demands large amounts of training data, and becomes inaccurate when generalized to unseen experiences. Remarkably, when LSTM or GRU are included in the neural network to retrofit the longitudinal command, these two types of RNN could also face difficulties in transfer learning, posing challenges in generalizing the model. Moreover, the stop-and-go speed profile and the fluctuated and choppy acceleration triggered by congestion appreciably contribute to the difficulties of retrofitting the vehicle longitudinal command, entailing a more intelligent neural network model to capture the fluctuations and discontinuities in the vehicle CF model during congestion. The poor data distribution generated from driver heterogeneity in congestion also contributes to the randomness of the model trained by BC, which casts extra doubt on the generalization of a BC model. Thus, BC could significantly suffer from the scarcity of training data and can be biased because of poor data distribution.

Although IRL and GAIL can circumvent some of the issues with the BC method, they still succumb to the pitfall of imitation learning methods. Chen et al. (111) summarized three major issues with imitation learning: (i) it needs to collect a huge amount of expert driving data in real-world and in real time, which can be costly and time-consuming, (ii) it can only learn driving skills that are demonstrated in the dataset. This might lead to serious issues given unseen experience in test process, and (iii) since the human driver experts act as the supervision for learning, it is impossible for an imitation learning policy to exceed human-level performance. From the traffic flow perspective, it is argued here that either BC or other deep imitation learning methods will be cumbersome, especially with incomplete datasets lacking the important driving scenarios mentioned above. According to Gao et al. (124), both BC and IRL algorithms implicitly assume that the demonstrations are complete, meaning that the action for each demonstrated state is fully observable and available. Obviously, this assumption does not hold for the mMP problem.

The existence of limitations with imitation learning methods highlights the potential of non-imitation

methods like RL in learning a “better” driving policy to reduce congestion and improve overall traffic efficiency. It is not easy to achieve, though. The major issue with using the RL method is the dependence on a reward function, which must be hand-crafted based on engineering experience and has to be applicable to all driving scenarios (125). RL methods might cause undesirable driving behaviors by directly transferring their driving policy learned in non-congestion states. Besides, it is argued here that adopting RL transforms the problem of mMP from imitating human demonstrations to searching for a policy that complies with a hand-crafted reward rule. Also, it should be pointed out that RL requires high-fidelity simulation platforms, which must be able to model accurately the appearance of the environment, the physics of vehicles, and the behavior of other participants (98). Especially important is the modeling of vehicle dynamics to represent the effects of gravity accurately, which has been found to be a key factor in reproducing empirical traffic flow instabilities (126).

In summary, RL seems to be the only hope to develop “optimal policies” that could potentially outperform human drivers. Despite the difficulty in designing a good reward function, and the requirement of a more realistic traffic environment, it is believed that the “trail-and-error” principle in RL is worth borrowing. Note that Tesla already seems to be working in this direction, and it is able to use the natural traffic environment to test its algorithms and collect ground-truth data. Again, it remains unknown whether Tesla has considered the congestion impact in its development program.

Traffic Domain Knowledge

Overall, the current research on mMP is devoting most of its efforts to the long tail safety problem, while its impact on congestion has been almost completely ignored. Through the above review, the major limitations in current datasets and learning methods have been identified, and now some potential future studies are proposed which aim at equipping the learning process with related traffic domain knowledge to fill in the research gap.

Here the main intellectual achievements in traditional CF theory are summarized which are probably worth noting for learning approaches to combine. Concerning the impact on traffic congestion, the most important human CF properties might include: memory and prediction, randomness, and string stability.

Memory and Prediction

For memory and prediction, LSTM, a type of RNN, has been adopted by mMP studies (30, 52, 127) to address the impact of memory on future speed choice. Lefèvre

et al. (123) conducted a comparative evaluation of parametric and non-parametric approaches for speed prediction during highway driving. Their study showed that the CF models can perform well for short-term speed prediction, but deep neural networks behave better for long-term prediction. To evaluate the relative performance of different learning methods on the same dataset, Kuefler et al. (108) compared the GAIL and BC methods using the same two-dimensional trajectories from NGSIM. Their work demonstrated that BC has the best short-horizon performance, and GAIL outperforms other methods including CF models for long-horizon tasks.

The CF models have realized the merits of introducing memory to improve prediction for a long time. Studies have also attempted to make some modifications to the traditional CF models based on their original form. Lee (128) revised the linear GHR model (129, 130) to account for the relative speed over a period of time: $a_n(t) = \int_0^t M(t-s)\Delta V_n(s)ds$, where M is the weight function for the memory impact. Tang et al. (131) extended the optimal velocity model (132) and found that considering human drivers' memory would improve the string stability of traffic flow. Similarly, Zhou et al. (90) captured traffic oscillations using a RNN-based CF model, which indicates that memory and prediction can help make informed driving decisions for smoother traffic.

It appears that mMP is able to imitate human driving with such memory and prediction property. For example, AVs will decelerate in advance when realizing potential decelerating or cut-in behaviors ahead of them. Notably, Elon Musk (98) also mentioned that Tesla can even predict a curving path that cannot be seen by humans because of road geometry or limited sight distance. The prediction power of mMP might outperform human behavior. Tesla has also demonstrated that its prediction can be used to infer the intention of other vehicles, such as cut-in behaviors which will be incorporated in AVs' motion planning. It is conjectured such prediction can improve traffic stability, because AVs can predict disruptive lane changes and prepare to decelerate first, instead of abrupt deceleration without any prediction. Those studies and new technologies pertinent to memory and prediction help to demonstrate the potential of AVs to dampen future traffic congestion.

String Stability and Safety

The literature has shown that most CF models are unable to replicate string stability consistent with empirical human driving data. These models are all deterministic, including stimulus-response models (130), optimal velocity models (133), IDM (Intelligent Driver Model) (134) and FVDM model (135), safe-distance model (136),

desired-headway model (137), and psycho-physical models (138, 139).

Sun et al. (140) conducted a comprehensive review on the methods for stability analysis and their applicability to CF models. They classified the traditional CF models into three categories: basic CF models, time-delayed CF models, and cooperative CF models, based on the assumption of a connected environment (140). Common methods in the literature for string stability analysis have also been reviewed in detail. However, those methods applicable for traditional CF models do not apply to mMP because of its lack of explicit mathematical formulations.

More importantly, Sun et al. pointed out some inconsistency between the results using analytical method and numerical simulation, which may result from some of the major assumptions or relaxations: (i) since the methods for string analysis are mostly based on linear equations, the non-linear CF models are approximated, which causes certain numerical errors; (ii) the platoon is always assumed to remain in equilibrium before a small perturbation is added when analyzing string stability, which goes against real traffic conditions where different driving regimes need to be considered, and (iii) the methods of linear stability analysis are only suitable for small perturbations and the non-linear effects caused by large perturbations such as hard braking do not apply. Those studies indicate that the string stability of mMP will be hard to capture because of the non-linear neural network architectures. Reasonable methods should depend on numerical studies. Therefore, to analyze the string stability of mMP, one has to approximate those proprietary "black boxes" with traditional CF models or a separate neural network, and then conduct numerical simulations for further analysis.

Moreover, safety (collision prevention) is another significant issue in mMP (actually it could be weighted the most in AV control design). In congested traffic, with the randomness and disturbances induced by human drivers, abrupt braking could be inevitable to guarantee safety, which could consequently jeopardize the string stability performance. Under this circumstance, how mMP will trade off collision avoidance and the smoothness of traffic in congestion remains to be analyzed and researched. A feasible direction could be making collision avoidance a local-level safety objective while using string stability as a system-level safety objective, and mMP will iteratively optimize these two objectives. Specifically, the local safety objective monitors the immediate safety status of the ego vehicle, preventing collisions with adjacent vehicles during driving tasks. The system-level safety objective could be evaluated as a long-term target, whose focus will be the smoothness (string stability) of the traffic. The reason is that the smooth traffic can alleviate the fluctuations of acceleration and enforce vehicles to operate closer to the

equilibrium, which can further prevent collisions in the surrounding traffic. Correspondingly, a specific boundary function needs to be scrutinizing the safety status during AV operation. Beyond the boundary, mMP can resort to optimizing the system-level performance to alleviate traffic oscillation, while within the boundary, the value of local safety will overwhelm the system-level string stability concern. Therefore, the smoothness of traffic can be an essential criteria of how AVs fit in the traffic in a long-term perspective, while collision prevention is the critical function for AVs to operate safely in a short time span. This is a significant issue to be carefully balanced such that AVs can scale up and benefit the traffic system.

Randomness

Laval et al. (141) showed that stochastic errors during the acceleration process are the cause of stop-and-go waves. They developed a parsimonious family of CF models that are able to reproduce most traffic instabilities, including traffic oscillations and capacity drop, based on stochastic processes to describe drivers' desired accelerations. It was found that this component is crucial for capturing realistic formation and propagation of traffic oscillations. This is probably the simplest CF model that captures driver random errors while accelerating and produces realistic traffic oscillations. Follow-up models that incorporate human error have also been formulated within this family (142, 143) and also for other well-known CF models (144).

To the best of the authors' knowledge, the stochastic property of mMP has not been well addressed or used for analyzing traffic congestion. It is not advisable, however, to add stochastic components to these methods because it will result in exacerbated traffic oscillations. On the contrary, one should try to minimize this error as much as possible, which should have a positive effect on congestion.

Connections Between CF Models and Neural Networks

While most mMP methods do not show a direct relationship with traditional CF models, it was revealed that a mathematical equivalence between mMP and CF models can be found under simple settings (145). A linear CF model will become interchangeable with a deep neural network given the same input and output. For equivalence in a real AV system, Xu et al. (52) showed that an mMP network can be replaced with a traditional CF model given speed and distance extracted from sensor data. It is argued here that mMP and CF models are mathematically equivalent if the mid-level methods generate position/distance-based learning affordances (features) as model input for mMP module. Since CF models adopt design variables of position and speed and

output acceleration, the mMP will boil down to a similar problem which maps the position or speed of surrounding cars to ego-vehicle acceleration. But such equivalence does not apply when the output of mMP becomes a predicted trajectory within a few seconds.

The mathematical connection between mMP and the CF models should result from the approximation power of neural networks, which has been discussed rigorously in the literature. Kolmogorov (146) proved a general theorem stating that any real-valued continuous function f defined on a n -dimension cube $I^n (n > 2)$ can be represented as:

$$f(x_1, x_2, \dots, x_n) = \sum_{q=1}^{2n+1} \phi_q \left(\sum_{p=1}^n \psi_{pq}(x_p) \right) \quad (3)$$

where ψ is a continuous and universal one-variable function, and ϕ is continuous monotonically increasing functions independent of f . Using Kolmogorov's theorem, Věra (147) also gave a direct proof of the universal approximation capabilities of perceptron networks with two hidden layers. Those studies may help to explain why neural networks can successfully replicate CF behaviors of human drivers and longitudinal control methods of AV.

Discussion and Outlook

This survey serves as a preliminary study to investigate the impact of AVs on traffic congestion in the future. It found that mMP is rapidly developing based on the efforts of the leading technology companies like Tesla and Comma.ai. Although mMP has not yet been widely applied, most automakers have already equipped enough hardware (sensors) to their latest car models to make mMP possible in the short-term future. Through the review it was also found that the AV industry has been mostly focusing on the long tail problem caused by "corner errors" related to safety, while the impact of AVs on traffic efficiency is almost ignored. In detail, none of the existing public datasets provides sufficient data that can be applied to the training of a congestion-mitigation mMP, and the major learning approach for mMP adopted by the industry is still BC. Albeit some non-imitation methods such as RL are proposed in the literature, there has not been noticeable success in training a congestion-mitigation or string-stable mMP for AVs in the existing literature, let alone its implementation in industry.

Research is needed to understand better the characteristics of mMP and their impact on traffic congestion. The authors suggest the following research directions.

Analyzing the Impact of AV by Approximation and Retrofitting

Since the current AV technologies are sealed as "black boxes," the only way to understand their behavior and

impact is to approximate and retrofit AVs using surrogate models. Noticing a certain level of equivalence between CF models and mMP, we can try to approximate the proprietary mMP by calibrating specific CF models. Similarly, in light of the universal approximation power of neural networks, it is also possible to find surrogate deep neural network models for currently unknown mMP models. Therefore, given a trained mMP, there are two different approaches to understanding its characteristics, either by calibrating a parameterized CF model or training a deep neural network as approximation. Both of the two methods will pave the way for further studies to analyze the impact of mMP on safety and string stability in traffic congestion.

Data Enrichment for Congestion-Oriented Research

Based on this investigation, there is insufficient data suitable for researching autonomous driving mMP in congested traffic. Most existing data are biased to emerging autonomous driving tasks such as object detection or safety issues in corner cases. Therefore, it is recommended that the industries and academic institutes should put more emphasis on the collection of data from AVs (not human-driven vehicles) in congestion, and potentially publish the data for further insights.

Incorporating Expert Knowledge from Traffic Domains

For future development of mMP it is advisable that planning agencies create incentives for the AV industry to put

more emphasis on the impact of AVs on traffic congestion, rather than only focusing on the long tail problem of “corner errors.” Relevant expert knowledge from traffic domains is worth noting, including but not limited to the properties of string stability revealed by traditional CF studies, impact of memory and prediction, the stochastic accelerations, and the equivalence between CF models and neural networks.

Conclusion

The paper has mainly surveyed and discussed the mMP for AVs, while leaving some other important factors including connectivity and the cooperation between AV industry and transportation agencies. The authors believe the emerging technology of connectivity also provides a great opportunity to benefit the traffic, as more real-time data enable AVs to execute traffic-friendly control algorithms. Additionally, the cooperation between AV industries and transportation agencies is also essential for improving the performance of AVs in congested traffic, and providing incentives for the smooth transition from human-driven vehicles to AVs.

Appendix. Major AV Technology Suppliers and Customers

A detailed graph showing the relations of major suppliers and customers in AV technology is included (Figure 5). A full table is shared via the link: <https://www.github.io/AV-technique-suppliers/>




Figure 5. Suppliers and customers of automated vehicle technology.
Source: Wu (29).

Author Contributions

The authors confirm contribution to the paper as follows: study conception and design: J. Laval, H. Zhou, S. Peeta; data collection: H. Zhou, W. Wu, A. Zhou, Y. Wang, Z. Qing; analysis and interpretation of results: H. Zhou, A. Zhou, Y. Wang, Z. Qing, W. Wu; draft manuscript preparation: H. Zhou, A. Zhou, J. Laval, Y. Wang, Z. Qing, W. Wu. All authors reviewed the results and approved the final version of the manuscript.

Declaration of Conflicting Interests

The author(s) declared no potential conflicts of interest with respect to the research, authorship, and/or publication of this article.

Funding

The author(s) disclosed receipt of the following financial support for the research, authorship, and/or publication of this article: This study is supported by NSF CPS grant \#1932451 and \#1826162.

ORCID iDs

Anye Zhou  <https://orcid.org/0000-0003-0145-5579>
 Yu Wang  <https://orcid.org/0000-0003-1054-555X>
 Zhu Qing  <https://orcid.org/0000-0002-3219-6971>
 Srinivas Peeta  <https://orcid.org/0000-0002-4146-6793>

References

1. Van Arem, B., C. J. Van Driel, and R. Visser. The Impact of Cooperative Adaptive Cruise Control on Traffic-Flow Characteristics. *IEEE Transactions on Intelligent Transportation Systems*, Vol. 7, No. 4, 2006, pp. 429–436.
2. Shladover, S., D. Su, and X.-Y. Lu. Impacts of Cooperative Adaptive Cruise Control on Freeway Traffic Flow. *Transportation Research Record: Journal of the Transportation Research Board*, No. 2324, 2012, pp. 63–70.
3. Talebpour, A., and H. S. Mahmassani. Influence of Connected and Autonomous Vehicles on Traffic Flow Stability and Throughput. *Transportation Research Part C: Emerging Technologies*, Vol. 71, 2016, pp. 143–163.
4. Mahmassani, H. S. 50th Anniversary Invited Article—Autonomous Vehicles and Connected Vehicle Systems: Flow and Operations Considerations. *Transportation Science*, Vol. 50, No. 4, 2016, pp. 1140–1162. doi: 10.1287/trsc.2016.0712.
5. Talebpour, A., H. S. Mahmassani, and S. H. Hamdar. Effect of Information Availability on Stability of Traffic Flow: Percolation Theory Approach. *Transportation Research Procedia*, Vol. 23, 2017, pp. 81–100. <http://dx.doi.org/10.1016/j.trpro.2017.05.006>.
6. Kesting, A., M. Treiber, M. Schönhof, and D. Helbing. Adaptive Cruise Control Design for Active Congestion Avoidance. *Transportation Research Part C: Emerging Technologies*, Vol. 16, No. 6, 2008, pp. 668–683.
7. Delis, A. I., I. K. Nikolos, and M. Papageorgiou. Macroscopic Traffic Flow Modeling with Adaptive Cruise Control: Development and Numerical Solution. *Computers & Mathematics with Applications*, Vol. 70, No. 8, 2015, pp. 1921–1947.
8. Gunter, G., C. Janssen, W. Barbour, R. E. Stern, and D. B. Work. Model Based String Stability of Adaptive Cruise Control Systems Using Field Data. *arXiv Preprint arXiv:1902.04983*, 2019.
9. Gunter, G., D. Gloudemans, R. E. Stern, S. McQuade, R. Bhadani, M. Bunting, M. L. Delle Monache, R. Lysecky, B. Seibold, J. Sprinkle, B. Piccoli, and D. B. Work. Are Commercially Implemented Adaptive Cruise Control Systems String Stable? *IEEE Transactions on Intelligent Transportation Systems*, 2020, pp. 1–12.
10. Li, T., D. Chen, H. Zhou, J. Laval, and Y. Xie. Car-Following Behavior Characteristics of Adaptive Cruise Control Vehicles Based on Empirical Experiments. *Transportation Research Part B: Methodological*, Vol. 147, 2021, pp. 67–91.
11. Naus, G. J., R. P. Vugts, J. Ploeg, M. J. van de Molengraft, and M. Steinbuch. String-Stable CACC Design and Experimental Validation: A Frequency-Domain Approach. *IEEE Transactions on vehicular technology*, Vol. 59, No. 9, 2010, pp. 4268–4279.
12. Naus, G., J. Ploeg, M. Van de Molengraft, W. Heemels, and M. Steinbuch. Design and Implementation of Parameterized Adaptive Cruise Control: An Explicit Model Predictive Control Approach. *Control Engineering Practice*, Vol. 18, No. 8, 2010, pp. 882–892. <https://doi.org/10.1016/j.conengprac.2010.03.012>.
13. Ploeg, J., B. T. M. Scheepers, E. van Nunen, N. van de Wouw, and H. Nijmeijer. Design and Experimental Evaluation of Cooperative Adaptive Cruise Control. *Proc., 14th International IEEE Conference on Intelligent Transportation Systems (ITSC)*, Washington, D.C., 2011, pp. 260–265.
14. Bu, F., H.-S. Tan, and J. Huang. Design and Field Testing of a Cooperative Adaptive Cruise Control System. *Proc., 2010 American Control Conference*, Baltimore, MD, 2010, pp. 4616–4621.
15. Zhou, A., S. Gong, C. Wang, and S. Peeta. Smooth-Switching Control-Based Cooperative Adaptive Cruise Control by Considering Dynamic Information Flow Topology. *Transportation Research Record*, Vol. 2674, 2020, pp. 444–458.
16. Zhou, H., A. Zhou, T. Li, D. Chen, S. Peeta, and J. Laval. Impact of the Low-Level Controller on StringStability of Adaptive Cruise Control System. *arXiv Preprint arXiv:2104.07726*, 2021.
17. Ajanovic, Z., B. Lacevic, B. Shyrokau, M. Stolz, and M. Horn. Search-Based Optimal Motion Planning for Automated Driving. *Proc., IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, IEEE, Madrid, Spain, 2018, pp. 4523–4530.
18. Tesla. Support. Autopilot and Full Self-Driving Capability. 2020. <https://www.tesla.com/support/autopilot> (accessed August 2, 2021).

19. Schmidt, B. Tesla to Ditch Radar in Coming Release of V9.0 Self-Driving Software. *The Driven*, June 8, 2020. <https://thedriven.io/2021/06/08/tesla-to-ditch-radar-in-coming-release-of-v9-0-self-driving-software/> (accessed August 2, 2021).
20. GMC Safety. Advanced Safety And Driver Assistance Features. *General Motors Corp.*, 2021. <https://www.gmc.com/safety-features>.
21. GM Authority. GM Super Cruise. 2020. <https://gmauthority.com/blog/gm/general-motors-technology/general-motors-autonomous-technology/gm-super-cruise/> (accessed August 2, 2021).
22. Radu, V. Nissan's ProPilot 2.0 Assist Tech: How It Works and What Are Its Limitations. 2020. <https://www.autoevolution.com/news/nissan-s-propilot-20-assist-how-it-works-and-what-are-its-limitations-152937.html>.
23. ProPILOT 2.0. Driver Assistance Technology, Combining Navigated Highway Driving with Hands-off Single-Lane Driving Capabilities. Nissan Motor Corp., 2020. <https://www.nissan-global.com/EN/TECHNOLOGY/OVERVIEW/ad2.html> (accessed August 2, 2021).
24. Toyota. Dynamic Radar Cruise Control (DRCC) or Full-Speed Range DRCC. 2020. <https://www.toyota.com/safety-sense/section/discover/feature/drcc/>.
25. Honda. Introduction to Honda Sensing. 2020. <https://global.honda/innovation/technology/automobile/Honda-Sensing.html>.
26. Ford. Introduction to Ford's IACC. 2020. <https://www.ford.com/technology/driver-assist-technology/adaptive-cruise-control/>.
27. Audi. Introduction to Audi's ACC with Stop-and-Go Function. 2020. <https://www.audi-technology-portal.de/en/electrics-electronics/driver-assistant-systems/adaptive-cruise-control-with-stop-go-function>.
28. BMW. Overview of the Main Driver Assistance System. 2020. <https://www.bmw.com/en/innovation/the-main-driver-assistance-systems.html>.
29. Wu, W. AV Technique Suppliers. 2020. <https://www20.github.io/AV-technique-suppliers/>.
30. Bansal, M., A. Krizhevsky, and A. Ogale. Chauffeurnet: Learning to Drive by Imitating the Best and Synthesizing the Worst. *arXiv Preprint arXiv:1812.03079*, 2018.
31. MOBILEYE 8 CONNECT, Driven by Safety: An AI-powered Fleet Safety Solution, 2020. <https://www.mobileye.com/us/fleets/products/mobileye-8-connect/>.
32. Scale. List of Public Self-Driving Datasets. 2020. <https://scale.com/open-datasets>.
33. Yu, H., S. Yang, W. Gu, and S. Zhang. Baidu Driving Dataset and End-to-End Reactive Control Model. *Proc., 2017 IEEE Intelligent Vehicles Symposium (IV)*, Los Angeles, CA, IEEE, 2017, pp. 341–346.
34. George, L., T. Buheit, E. Wirbel, G. Le-Gall, and X. Perrotton. Imitation Learning for End to End Vehicle Longitudinal Control with Forward Camera. *arXiv Preprint arXiv:1812.05841*, 2018.
35. Sharma, S., G. Tewolde, and J. Kwon. Lateral and Longitudinal Motion Control of Autonomous Vehicles Using Deep Learning. *IEEE International Conference on Electro Information Technology*, Vol. May, 2019, pp. 460–464. <https://doi.org/10.1109/EIT.2019.8833873>.
36. Kuutti, S., R. Bowden, H. Joshi, R. de Temple, and S. Fallah. End-to-End Reinforcement Learning for Autonomous Longitudinal Control Using Advantage Actor Critic with Temporal Context. *Proc., 2019 IEEE Intelligent Transportation Systems Conference (ITSC)*, Auckland, New Zealand, and, IEEE, 2019, pp. 2456–2462.
37. Pathak, S., V. J. Nadkarni, and S. Bag. Adaptive Longitudinal Control Using Reinforcement Learning. U.S. Patent App. 16/427,589, 2019.
38. Babak, S.-J., S. A. Hussain, B. Karakas, and S. Cetin. Control of Autonomous Ground Vehicles: A Brief Technical Review. *IOP Conference Series: Materials Science and Engineering*, Vol. 224, 2017, p. 012029.
39. Katrakazas, C., M. Quddus, W.-H. Chen, and L. Deka. Real-Time Motion Planning Methods for Autonomous On-Road Driving: State-of-the-Art and Future Research Directions. *Transportation Research Part C: Emerging Technologies*, Vol. 60, 2015, pp. 416–442.
40. Paden, B., M. Čáp, S. Z. Yong, D. Yershov, and E. Frazzoli. A Survey of Motion Planning and Control Techniques for Self-Driving Urban Vehicles. *IEEE Transactions on Intelligent Vehicles*, Vol. 1, No. 1, 2016, pp. 33–55.
41. Dey, K. C., L. Yan, X. Wang, Y. Wang, H. Shen, M. Chowdhury, L. Yu, C. Qiu, and V. Soundararaj. A Review of Communication, Driver Characteristics, and Controls Aspects of Cooperative Adaptive Cruise Control (CACC). *IEEE Transactions on Intelligent Transportation Systems*, Vol. 17, No. 2, 2015, pp. 491–509.
42. Ni, J., Y. Chen, Y. Chen, J. Zhu, D. Ali, and W. Cao. A Survey on Theories and Applications for Self-Driving Cars Based on Deep Learning Methods. *Applied Sciences*, Vol. 10, No. 8, 2020, p. 2749.
43. Schwarting, W., J. Alonso-Mora, and D. Rus. Planning and Decision-Making for Autonomous Vehicles. *Annual Review of Control, Robotics, and Autonomous Systems*, Vol. 1, 2018, pp. 187–210.
44. Yurtsever, E., J. Lambert, A. Carballo, and K. Takeda. A Survey of Autonomous Driving: Common Practices and Emerging Technologies. *arXiv Preprint arXiv:1906.05113*, 2019.
45. Di, X., and R. Shi. A Survey on Autonomous Vehicle Control in the Era of Mixed-Autonomy: From Physics-Based to AI-Guided Driving Policy Learning. *arXiv Preprint arXiv:2007.05156*, 2020.
46. Talpaert, V., I. Sobh, B. R. Kiran, P. Mannion, S. Yogamani, A. El-Sallab, and P. Perez. Exploring Applications of Deep Reinforcement Learning for Real-World Autonomous Driving Systems. *arXiv Preprint arXiv:1901.01536*, 2019.
47. Yin, H., and C. Berger. When to Use What Data Set for Your Self-Driving Car Algorithm: An Overview of Publicly Available Driving Datasets. *Proc., IEEE 20th International Conference on Intelligent Transportation Systems (ITSC)*, Yokohama, Japan, IEEE, 2017, pp. 1–8.
48. Kang, Y., H. Yin, and C. Berger. Test Your Self-Driving Algorithm: An Overview of Publicly Available Driving Datasets and Virtual Testing Environments. *IEEE*

- Transactions on Intelligent Vehicles*, Vol. 4, No. 2, 2019, pp. 171–185.
49. Sun, P., H. Kretzschmar, X. Dotiwalla, A. Chouard, V. Patnaik, P. Tsui, J. Guo, Y. Zhou, Y. Chai, B. Caine, V. Vasudevan, W. Han, J. Ngiam, H. Zhao, A. Timofeev, S. Ettinger, M. Krivokon, A. Gao, A. Joshi, Y. Zhang, J. Shlens, Z. Chen, and D. Anguelov. Scalability in Perception for Autonomous Driving: Waymo Open Dataset. *arXiv Preprint arXiv:1912.04838*, 2019.
 50. Huang, X., X. Cheng, Q. Geng, B. Cao, D. Zhou, P. Wang, Y. Lin, and R. Yang. The Apolloscape Dataset for Autonomous Driving. *Proc., IEEE Conference on Computer Vision and Pattern Recognition Workshops*. 2018, pp. 954–960.
 51. Geiger, A., P. Lenz, and R. Urtasun. Are We Ready for Autonomous Driving? The KITTI Vision Benchmark Suite. *Proc., 2012 IEEE Conference on Computer Vision and Pattern Recognition*, Providence, RI, IEEE, 2012, pp. 3354–3361.
 52. Xu, H., Y. Gao, F. Yu, and T. Darrell. End-to-End Learning of Driving Models from Large-Scale Video Datasets. *Proc., 30th IEEE Conference on Computer Vision and Pattern Recognition*, Los Alamitos, CA, 2017, pp. 3530–3538.
 53. Forson, E. Teaching Cars to Drive—Highway Path Planning. *Towards Data Science*, July 8, 2018. <https://towardsdatascience.com/teaching-cars-to-drive-highway-path-planning-109c49f9f86c> (accessed August 2, 2021).
 54. Caesar, H., V. Bankiti, A. H. Lang, S. Vora, V. E. Liong, Q. Xu, A. Krishnan, Y. Pan, G. Baldan, and O. Beijbom. nuScenes: A Multimodal Dataset for Autonomous Driving. *arXiv Preprint arXiv:1903.11027*, 2019.
 55. Agarwal, S., A. Vora, G. Pandey, W. Williams, H. Kourous, and J. R. McBride. Ford Multi-AV Seasonal Dataset. *arXiv Preprint arXiv:2003.07969*, 2020.
 56. Chang, M.-F., J. W. Lambert, P. Sangkloy, J. Singh, S. Bak, A. Hartnett, D. Wang, P. Carr, S. Lucey, D. Ramanan, and J. Hays. Argoverse: 3D Tracking and Forecasting with Rich Maps. *Proc., Conference on Computer Vision and Pattern Recognition*, Long Beach, CA, 2019, pp. 8740–8749.
 57. Schafer, H., E. Santana, A. Haden, and R. Biasini. A Commute in Data: The comma2k19 Dataset. *arXiv Preprint arXiv:1812.05752*, 2018.
 58. Jain, A., H. S. Koppula, S. Soh, B. Raghavan, A. Singh, and A. Saxena. Brain4Cars: Car That Knows Before You Do via Sensory-Fusion Deep Learning Architecture. *arXiv Preprint arXiv: 1601.00740*, 2016.
 59. Cordts, M., M. Omran, S. Ramos, T. Rehfeld, M. Enzweiler, R. Benenson, U. Franke, S. Roth, and B. Schiele. The Cityscapes Dataset for Semantic Urban Scene Understanding. *Proc., IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, Las Vegas, NV, 2016.
 60. Barnes, D., M. Gadd, P. D. Murcutt, P. Newman, and I. Posner. The Oxford Radar RobotCar Dataset: A Radar Extension to the Oxford RobotCar Dataset. *arXiv Preprint arXiv: 1909.01300*, 2019.
 61. Romera, E., L. M. Bergasa, and R. Arroyo. Need Data for Driver Behaviour Analysis? Presenting the Public UAH-DriveSet. *Proc., IEEE 19th International Conference on Intelligent Transportation Systems (ITSC)*, Rio de Janeiro, Brazil, 2016, pp. 387–392.
 62. Krajewski, R., J. Bock, L. Kloeker, and L. Eckstein. The highD Dataset: A Drone Dataset of Naturalistic Vehicle Trajectories on German Highways for Validation of Highly Automated Driving Systems. *Proc., 21st International Conference on Intelligent Transportation Systems (ITSC)*, Maui, HI, 2018, pp. 2118–2125.
 63. Hiller, J., S. Koskinen, R. Berta, N. Osman, B. Nagy, F. Bellotti, A. Rahman, E. Svanberg, H. Weber, E. Arnold, M. Dianati, and A. D. Gloria. The L3Pilot Data Management Toolchain for a Level 3 Vehicle Automation Pilot. *Electronics*, Vol. 9, 2020, p. 809.
 64. Makridis, M., K. Mattas, A. Anesiadou, and B. Ciuffo. openACC. An Open Database of Car-Following Experiments to Study the Properties of Commercial ACC Systems. *arXiv Preprint arXiv:2004.06342*, 2020.
 65. Kesten, R., M. Usman, J. Houston, T. Pandya, K. Nadhamuni, A. Ferreira, M. Yuan, B. Low, A. Jain, P. Ondruska, S. Omari, S. Shah, A. Kulkarni, A. Kazakova, C. Tao, L. Platinsky, W. Jiang, and V. Shet. Lyft Level 5 AV Dataset. 2019. <https://level5.lyft.com/dataset/>.
 66. Dosovitskiy, A., G. Ros, F. Codevilla, A. Lopez, and V. Koltun. CARLA: An Open Urban Driving Simulator. *arXiv Preprint arXiv:1711.03938*, 2017.
 67. Wymann, B., E. Espié, C. Guionneau, C. Dimitrakakis, R. Coulom, and A. Sumner. Torcs, the Open Racing Car Simulator. *Software Available at http://torcs.sourceforge.net*, Vol. 4, No. 6, 2000, p. 2.
 68. Chen, C., A. Seff, A. Kornhauser, and J. Xiao. Deepdriving: Learning Affordance for Direct Perception in Autonomous Driving. *Proc., IEEE International Conference on Computer Vision*, Santiago, Chile, 2015, pp. 2722–2730.
 69. Panwai, S., and H. Dia. Neural Agent Car-Following Models. *IEEE Transactions on Intelligent Transportation Systems*, Vol. 8, No. 1, 2007, pp. 60–70.
 70. Codevilla, F., M. Müller, A. López, V. Koltun, and A. Dosovitskiy. End-to-End Driving Via Conditional Imitation Learning. *Proc., 2018 IEEE International Conference on Robotics and Automation (ICRA)*, Brisbane, QLD, Australia, IEEE, 2018, pp. 1–9.
 71. Mirowski, P., M. Grimes, M. Malinowski, K. M. Hermann, K. Anderson, D. Teplyashin, K. Simonyan, A. Zisserman, K. Kavukcuoglu, and R. Hadsell. Learning to Navigate in Cities Without a Map. *Proc., 32nd Conference on Advances in Neural Information Processing Systems*, Montréal, Canada, 2018, pp. 2419–2430.
 72. Tan, B., N. Xu, and B. Kong. Autonomous Driving in Reality with Reinforcement Learning and Image Translation. *arXiv Preprint arXiv:1801.05299*, 2018.
 73. Müller, M., A. Dosovitskiy, B. Ghanem, and V. Koltun. Driving Policy Transfer via Modularity and Abstraction. *arXiv Preprint arXiv:1804.09364*, 2018.
 74. Wei, J., and J. M. Dolan. A Robust Autonomous Freeway Driving Algorithm. *Proc., 2009 IEEE Intelligent Vehicles Symposium*, IEEE, Xi'an, China, 2009, pp. 1015–1020.
 75. Leonard, J., J. How, S. Teller, M. Berger, S. Campbell, G. Fiore, L. Fletcher, E. Frazzoli, A. Huang, S. Karaman, O.

- Koch, Y. Kuwata, D. Moore, E. Olson, S. Peters, J. Teo, R. Truax, and M. Walter. A Perception-Driven Autonomous Urban Vehicle. *Journal of Field Robotics*, Vol. 25, No. 10, 2008, pp. 727–774.
76. Dabir, L., C. Goodin, N. Scherrer, and D. Carruth. LiDAR Data Segmentation in Off-Road Environment Using Convolutional Neural Networks (CNN). *SAE International Journal of Advances and Current Practices in Mobility*, Vol. 2, No. 6, 2020, pp. 3288–3292.
77. Feng, S., X. Yan, H. Sun, Y. Feng, and H. X. Liu. Intelligent Driving Intelligence Test for Autonomous Vehicles with Naturalistic and Adversarial Environment. *Nature Communications*, Vol. 12, No. 1, 2021, pp. 1–14.
78. Short, M., and M. J. Pont. Assessment of High-Integrity Embedded Automotive Control Systems Using Hardware in the Loop Simulation. *Journal of Systems and Software*, Vol. 81, No. 7, 2008, pp. 1163–1183.
79. Zhao, J., H. Wu, and C. Chang. *Virtual Traffic Simulator for Connected and Automated Vehicles*. SAE Technical Paper 2019-01-0676. Society of Automotive Engineers, Detroit, 2019.
80. Cantas, M. R., S. Fan, O. Kavas, S. Tamilarasan, L. Guvenc, S. Yoo, J. H. Lee, B. Lee, and J. Ha. *Development of Virtual Fuel Economy Trend Evaluation Process*. SAE Technical Paper 2019-01-0510. Society of Automotive Engineers, Detroit, MI, 2019.
81. Kehrer, M., J. Pitz, T. Rothermel, and H.-C. Reuss. Framework for Interactive Testing and Development of Highly Automated Driving Functions. In 18. Internationales Stuttgarter Symposium (M. Bargende, H. C. Reuss, and J. Wiedemann, eds.), Springer, Wiesbaden, 2018, pp. 659–669.
82. Wiggers, K. Uber Open-Sources Autonomous Visualization System, a Web-Based Platform for Vehicle Data. *VB*, February 19, 2019. <https://venturebeat.com/2019/02/19/uber-open-sources-autonomous-visualization-system-a-web-based-platform-for-vehicle-data/> (accessed August 2, 2021).
83. Wu, C., A. Kreidieh, K. Parvate, E. Vinitsky, and A. M. Bayen. Flow: Architecture and Benchmarking for Reinforcement Learning in Traffic Control. *arXiv Preprint arXiv:1710.05465*, 2017.
84. Kreidieh, A. R., C. Wu, and A. M. Bayen. Dissipating Stop-and-Go Waves in Closed and Open Networks Via Deep Reinforcement Learning. *Proc., 21st International Conference on Intelligent Transportation Systems (ITSC)*, Maui, HI, IEEE, 2018, pp. 1475–1480.
85. Kim, J., and J. Canny. Interpretable Learning for Self-Driving Cars by Visualizing Causal Attention. *Proc., IEEE International Conference on Computer Vision*, Venice, Italy, 2017, pp. 2942–2950.
86. Bojarski, M., D. Del Testa, D. Dworakowski, B. Firner, B. Flepp, P. Goyal, L. D. Jackel, M. Monfort, U. Muller, J. Zhang, X. Zhang, J. Zhao, and K. Zieba. End to End Learning for Self-Driving Cars. *arXiv Preprint arXiv:1604.07316*, 2016.
87. Chen, Z., and X. Huang. End-to-End Learning for Lane Keeping of Self-Driving Cars. *Proc., IEEE Intelligent Vehicles Symposium*, Los Angeles, CA, No. Iv, 2017, pp. 1856–1860.
88. Eraqi, H. M., M. N. Moustafa, and J. Honer. End-to-End Deep Learning for Steering Autonomous Vehicles Considering Temporal Dependencies. *arXiv Preprint arXiv:1710.03804*, 2017.
89. Hecker, S., D. Dai, and L. Van Gool. End-to-End Learning of Driving Models with Surround-View Cameras and Route Planners. *Proc., European Conference on Computer Vision (ECCV)*, Munich, Germany, 2018, pp. 435–453.
90. Zhou, M., X. Qu, and X. Li. A Recurrent Neural Network Based Microscopic Car Following Model to Predict Traffic Oscillation. *Transportation Research Part C: Emerging Technologies*, Vol. 84, 2017, pp. 245–264. <http://dx.doi.org/10.1016/j.trc.2017.08.027>.
91. Deo, N., and M. M. Trivedi. Convolutional Social Pooling for Vehicle Trajectory Prediction. *IEEE Computer Society Conference on Computer Vision and Pattern Recognition Workshops*, Vol. June, 2018, pp. 1549–1557. <https://doi.org/10.1109/CVPRW.2018.00196.1805.06771>.
92. Lee, D., Y. Gu, J. Hoang, and M. Marchetti-Bowick. Joint Interaction and Trajectory Prediction for Autonomous Driving Using Graph Neural Networks. *arXiv Preprint arXiv:1912.07882*, 2019.
93. Su, J., P. A. Beling, R. Guo, and K. Han. Graph Convolution Networks for Probabilistic Modeling of Driving Acceleration. *arXiv Preprint arXiv:1911.09837*, 2019.
94. Jeon, H., J. Choi, and D. Kum. SCALE-Net: Scalable Vehicle Trajectory Prediction Network under Random Number of Interacting Vehicles via Edge-Enhanced Graph Convolutional Neural Network. *arXiv Preprint arXiv:2002.12609*, 2020.
95. Yang, Z., Y. Zhang, J. Yu, J. Cai, and J. Luo. End-to-End Multi-Modal Multi-Task Vehicle Control for Self-Driving Cars with Visual Perceptions. *Proc., 24th International Conference on Pattern Recognition*, August, 2018, pp. 2289–2294. <https://doi.org/10.1109/ICPR.2018.8546189>.1801.06734.
96. Hsu, T. M., C. H. Wang, and Y. R. Chen. End-to-End Deep Learning for Autonomous Longitudinal and Lateral Control Based on Vehicle Dynamics. *Proc., ACM International Conference Proceeding Series*, Nagoya, Japan, 2018, pp. 111–114.
97. Li, Z., Z. Gu, X. Di, and R. Shi. An LSTM-Based Autonomous Driving Model Using Waymo Open Dataset. *arXiv Preprint arXiv:2002.05878*, 2020.
98. Musk, E. Tesla Autonomy Investor Day. 2019. https://www.youtube.com/watch?v=Ucp0TTmvqOE&ab_channel=Tesla.
99. Future of Driving, Tesla, 2020. <https://www.tesla.com/autopilot>. (accessed August 2, 2021)
100. Karpathy, A., AI for Full Self-Driving. 5th Annual ScaledML 2020 Conference, February 26–27, 2020. <https://info.matroid.com/scaledml-media-archive-2020> (accessed August 2, 2021).
101. Finn, C., P. Christiano, P. Abbeel, and S. Levine. A Connection Between Generative Adversarial Networks, Inverse Reinforcement Learning, and Energy-Based Models. *arXiv Preprint arXiv:1611.03852*, 2019.

102. Abbeel, P., and A. Y. Ng. Apprenticeship Learning via Inverse Reinforcement Learning. *Proc., 21st International Conference on Machine Learning*, Banff, Alberta, Canada, ACM, 2004, p. 1.
103. Sadigh, D., S. Sastry, S. A. Seshia, and A. D. Dragan. Planning for Autonomous Cars that Leverage Effects on Human Actions. *Robotics: Science and Systems*, Vol. 2, Ann Arbor, MI, 2016.
104. González, D. S., J. S. Dibangoye, and C. Laugier. High-Speed Highway Scene Prediction Based on Driver Models Learned from Demonstrations. *Proc., 2016 IEEE 19th International Conference on Intelligent Transportation Systems (ITSC)*, Rio de Janeiro, Brazil, IEEE, 2016, pp. 149–155.
105. Sharifzadeh, S., I. Chiotellis, R. Triebel, and D. Cremers. Learning to Drive Using Inverse Reinforcement Learning and Deep Q-Networks. *arXiv Preprint arXiv:1612.03653*, 2016.
106. Ziebart, B. D., A. Maas, J. A. Bagnell, and A. K. Dey. Maximum Entropy Inverse Reinforcement Learning. *arXiv Preprint arXiv:1507.04888*, 2015.
107. Ho, J., and S. Ermon. Generative Adversarial Imitation Learning. *Advances in Neural Information Processing Systems*, Vol. 29, 2016, pp. 4565–4573.
108. Kuefler, A., J. Morton, T. Wheeler, and M. Kochenderfer. Imitating Driver Behavior with Generative Adversarial Networks. *Proc., 2017 IEEE Intelligent Vehicles Symposium (IV)*, Los Angeles, CA, IEEE, 2017, pp. 204–211.
109. Bhattacharyya, R. P., D. J. Phillips, B. Wulfe, J. Morton, A. Kuefler, and M. J. Kochenderfer. Multi-Agent Imitation Learning for Driving Simulation. *Proc., 2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, Madrid, Spain, IEEE, 2018, pp. 1534–1539.
110. Pan, X., Y. You, Z. Wang, and C. Lu. Virtual to Real Reinforcement Learning for Autonomous Driving. *arXiv Preprint arXiv:1704.03952*, 2017.
111. Chen, J., B. Yuan, and M. Tomizuka. Model-Free Deep Reinforcement Learning for Urban Autonomous Driving. *arXiv Preprint arXiv:1904.09503*, 2019.
112. Guo, Q., O. Angah, Z. Liu, and X. J. Ban. Hybrid Deep Reinforcement Learning Based Eco-Driving for Low-Level Connected and Automated Vehicles along Signalized Corridors. *Transportation Research Part C: Emerging Technologies*, Vol. 124, 2021, p. 102980.
113. Shalev-Shwartz, S., S. Shammah, and A. Shashua. Safe, Multi-Agent, Reinforcement Learning for Autonomous Driving. *arXiv Preprint arXiv:1610.03295*, 2016.
114. Fridman, L. DeepTraffic: Crowdsourced Hyperparameter Tuning of Deep Reinforcement Learning Systems for Multi-Agent Dense Traffic Navigation. In *Neural Information Processing Systems (NIPS 2018) Deep Reinforcement Learning Workshop*, Montréal, CA, 2018.
115. Sallab, A. E., M. Abdou, E. Perot, and S. Yogamani. Deep Reinforcement Learning Framework for Autonomous Driving. *Electronic Imaging*, Vol. 2017, No. 19, 2017, pp. 70–76.
116. Kendall, A., J. Hawke, D. Janz, P. Mazur, D. Reda, J.-M. Allen, V.-D. Lam, A. Bewley, and A. Shah. Learning to Drive in a Day. *arXiv Preprint arXiv:1807.00412*, 2018.
117. Liang, X., T. Wang, L. Yang, and E. Xing. CIRC: Controllable Imitative Reinforcement Learning for Vision-Based Self-Driving. *Proc., European Conference on Computer Vision (ECCV)*, Munich, Germany, 2018, pp. 584–599.
118. Stern, R. E., S. Cui, M. L. Delle Monache, R. Bhadani, M. Bunting, M. Churchill, N. Hamilton, H. Pohlmann, F. Wu, B. Piccoli, B. Seibold, J. Sprinkle, and D.B. Work. Dissipation of Stop-and-Go Waves Via Control of Autonomous Vehicles: Field Experiments. *Transportation Research Part C: Emerging Technologies*, Vol. 89, 2018, pp. 205–221.
119. Codevilla, F., E. Santana, A. M. López, and A. Gaidon. Exploring the Limitations of Behavior Cloning for Autonomous Driving. *arXiv Preprint arXiv:1904.08980*, 2019.
120. Laval, J. A., and L. Leclercq. Microscopic Modeling of the Relaxation Phenomenon Using a Macroscopic Lane-Changing Model. *Transportation Research Part B: Methodological*, Vol. 42, No. 6, 2008, pp. 511–522.
121. Sauer, A., N. Savinov, and A. Geiger. Conditional Affordance Learning for Driving in Urban Environments. *arXiv Preprint arXiv:1806.06498*, 2018.
122. Wheeler, T. A., P. Robbel, and M. J. Kochenderfer. Analysis of Microscopic Behavior Models for Probabilistic Modeling of Driver Behavior. *Proc., IEEE 19th International Conference on Intelligent Transportation Systems (ITSC)*, Rio de Janeiro, Brazil, IEEE, 2016, pp. 1604–1609.
123. Lefèvre, S., C. Sun, R. Bajcsy, and C. Laugier. Comparison of Parametric and Non-Parametric Approaches for Vehicle Speed Prediction. *Proc., 2014 American Control Conference*, Portland, OR, IEEE, 2014, pp. 3494–3499.
124. Gao, Y., J. Lin, F. Yu, S. Levine, and T. Darrell. Reinforcement Learning from Imperfect Demonstrations. *arXiv Preprint arXiv:1802.05313*, 2018.
125. Makantasis, K., M. Kontorinaki, and I. Nikolos. A Deep Reinforcement Learning Driving Policy for Autonomous Road Vehicles. *arXiv Preprint arXiv:1905.09046*, 2019.
126. Laval, J. A. *Hybrid Models of Traffic Flow: Impacts of Bounded Vehicle Accelerations*. University of California, Berkeley, CA, 2004.
127. Morton, J., and M. J. Kochenderfer. Simultaneous Policy Learning and Latent State Inference for Imitating Driver Behavior. *Proc., IEEE 20th International Conference on Intelligent Transportation Systems (ITSC)*, Yokohama, Japan, IEEE, 2017, pp. 1–6.
128. Lee, G. A Generalization of Linear Car-Following Theory. *Operations Research*, Vol. 14, No. 4, 1966, pp. 595–606.
129. Chandler, R. E., R. Herman, and E. W. Montroll. Traffic Dynamics: Studies in Car Following. *Operations Research*, Vol. 6, No. 2, 1958, pp. 165–184. <https://doi.org/10.1287/opre.6.2.165>.
130. Gazis, D. C., R. Herman, and R. B. Potts. Car-Following Theory of Steady-State Traffic Flow. *Operations Research*, Vol. 7, No. 4, 1959, pp. 499–505.

131. Tang, T., H. Huang, S. Zhao, and G. Xu. An Extended OV Model with Consideration of Driver's Memory. *International Journal of Modern Physics B*, Vol. 23, No. 05, 2009, pp. 743–752.
132. Bando, M., K. Hasebe, K. Nakanishi, and A. Nakayama. Analysis of Optimal Velocity Model with Explicit Delay. *Physical Review E*, Vol. 58, No. 5, 1998, p. 5429.
133. Bando, M., K. Hasebe, A. Nakayama, A. Shibata, and Y. Sugiyama. Dynamical Model of Traffic Congestion and Numerical Simulation. *Physical Review E*, Vol. 51, No. 2, 1995, p. 1035.
134. Treiber, M., A. Hennecke, and D. Helbing. Congested Traffic States in Empirical Observations and Microscopic Simulations. *Physical Review E*, Vol. 62, No. 2, 2000, p. 1805.
135. Jiang, R., Q. Wu, and Z. Zhu. Full Velocity Difference Model for a Car-Following Theory. *Physical Review E*, Vol. 64, No. 1, 2001, p. 017101.
136. Gipps, P. G. A Behavioural Car-Following Model for Computer Simulation. *Transportation Research Part B: Methodological*, Vol. 15, No. 2, 1981, pp. 105–111.
137. Bullen, A. Development of Compact Microsimulation for Analyzing Freeway Operations and Design. Presented at 61st Annual Meeting of the Transportation Research Board, Washington, D.C., 1982.
138. Michaels, R. Perceptual Factors in Car-Following. *Proc., 2nd International Symposium on the Theory of Road Traffic Flow*, OECD, Paris, 1963, pp. 44–59.
139. Wiedemann, R. *Simulation des Strassenverkehrsflusses*. Institut für Verkehrswesen der Universität Karlsruhe, Germany, 1974.
140. Sun, J., Z. Zheng, and J. Sun. Stability Analysis Methods and Their Applicability to Car-Following Models in Conventional and Connected Environments. *Transportation Research Part B: Methodological*, Vol. 109, 2018, pp. 212–237.
141. Laval, J. A., C. S. Toth, and Y. Zhou. A Parsimonious Model for the Formation of Oscillations in Car-Following Models. *Transportation Research Part B: Methodological*, Vol. 70, 2014, pp. 228–238. <https://doi.org/10.1016/j.trb.2014.09.004>.
142. Xu, T., and J. Laval. Analysis of a Two-Regime Stochastic Car-Following Model: Explaining Capacity Drop and Oscillation Instabilities. *Transportation Research Record: Journal of the Transportation Research Board*, 2019. 2673: 610–619.
143. Yuan, K., J. Laval, V. L. Knoop, R. Jiang, and S. Hoogendoorn. A Geometric Brownian Motion Car-Following Model: Towards a Better Understanding of Capacity Drop. *Transportmetrica B: Transport Dynamics*, Vol. 7, No. 1, 2019, pp. 915–927.
144. Treiber, M., and A. Kesting. The Intelligent Driver Model with Stochasticity - New Insights into Traffic Flow Oscillations. *Transportation Research Procedia*, Vol. 23, No. Supplement C, 2017, pp. 174–187. <https://doi.org/10.1016/j.trpro.2017.05.011>.
145. Wu, F., and D. B. Work. Connections Between Classical Car Following Models and Artificial Neural Networks. *Proc., 21st International Conference on Intelligent Transportation Systems (ITSC)*, Maui, HI, IEEE, 2018, pp. 3191–3198.
146. Kolmogorov, A. N. On the Representation of Continuous Functions of Many Variables by Superposition of Continuous Functions of One Variable and Addition. *Doklady Akademii Nauk*, Vol. 114, No. 5, 1957, pp. 953–956.
147. Věra, K. Kolmogorov's Theorem and Multilayer Neural Networks. *Neural Networks*, Vol. 5, No. 3, 1992, pp. 501–506.