

Artificial Intelligence for Automated Vehicle Control and Traffic Operations Challenges and Opportunities

Abbink, David A.; Hao, Peng; Laval, Jorge; Shalev-Shwartz, Shai; Wu, Cathy; Yang, Terry; Hamdar, Samer; Chen, Danjue; Xie, Yuanchang; Li, Xiaopeng

DOI

[10.1007/978-3-030-80063-5_6](https://doi.org/10.1007/978-3-030-80063-5_6)

Publication date

2022

Document Version

Final published version

Published in

Road Vehicle Automation 8

Citation (APA)

Abbink, D. A., Hao, P., Laval, J., Shalev-Shwartz, S., Wu, C., Yang, T., Hamdar, S., Chen, D., Xie, Y., Li, X., & Haque, M. (2022). Artificial Intelligence for Automated Vehicle Control and Traffic Operations: Challenges and Opportunities. In G. Meyer, & S. Beiker (Eds.), *Road Vehicle Automation 8* (pp. 60-72). (Lecture Notes in Mobility). Springer. https://doi.org/10.1007/978-3-030-80063-5_6

Important note

To cite this publication, please use the final published version (if applicable).
Please check the document version above.

Copyright

Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

Takedown policy

Please contact us and provide details if you believe this document breaches copyrights.
We will remove access to the work immediately and investigate your claim.

Green Open Access added to TU Delft Institutional Repository

'You share, we take care!' - Taverne project

<https://www.openaccess.nl/en/you-share-we-take-care>

Otherwise as indicated in the copyright section: the publisher is the copyright holder of this work and the author uses the Dutch legislation to make this work public.



Artificial Intelligence for Automated Vehicle Control and Traffic Operations: Challenges and Opportunities

David A. Abbink¹, Peng Hao², Jorge Laval³, Shai Shalev-Shwartz⁴, Cathy Wu⁵, Terry Yang⁶, Samer Hamdar⁷(✉), Danjue Chen⁸, Yuanchang Xie⁹, Xiaopeng Li¹⁰, and Mohaiminul Haque⁷

¹ Section Human-Robot Interaction, Department of Cognitive Robotics, Faculty 3mE, Delft University of Technology, Room: 34-F-2-120, Mekelweg 2, 2628, CD Delft, The Netherlands

D.A.Abbink@tudelft.nl

² University of California Riverside, Riverside, USA

peng.hao@ucr.edu

³ Georgia Institute of Technology, SEB 224, 788 Atlantic Drive NW, 30332 Atlanta, GA, Georgia

jorge.laval@ce.gatech.edu

⁴ Mobileye and The Hebrew University of Jerusalem, Rothberg Building, Room B428, Givat Ram, 91904 Jerusalem, Israel

shais@cs.huji.ac.il

⁵ Massachusetts Institute of Technology, 77 Massachusetts Avenue, Cambridge, MA 02139-4307, USA

cathywu@mit.edu

⁶ The University of Utah, Floyd and Jeri Meldrum Civil Engineering, 110 Central Campus Drive, UT 84112 Salt Lake City, USA

x.yang@utah.edu

⁷ George Washington University, 800 22nd Street NW, Washington, DC 20052, USA

{hamdar,mohaiminul}@gwu.edu

⁸ University of Massachusetts Lowell, 1 University Avenue, PA 108, MA 01854 Lowell, USA

danjue_chen@uml.edu

⁹ University of Massachusetts Lowell, 1 University Avenue, MA 01854 Lowell, USA

Yuanchang_Xie@uml.edu

¹⁰ University of South Florida, 4202 E. Fowler Avenue, ENG 207, Tampa, FL 33620-5350, USA

xiaopengli@usf.edu

Abstract. This chapter summarizes the presentations of speakers addressing such issues during the Automated Vehicles Symposium 2020 (AVS20) held virtually on July 27–30, 2020. These speakers participated in the break-out session titled “Artificial Intelligence for Automated Vehicle Control and Traffic Operations: Challenges and Opportunities”. The corresponding discussion and recommendations are presented in terms of the lessons learned and the future research directions to be adopted to benefit from AI in order to develop safer and more efficient connected and automated vehicles (CAV). This session was organized by the Transportation Research Board (TRB) Committee on Traffic Flow Theory and Characteristics (ACP50) and the TRB Committee on Artificial Intelligence and Advanced Computing Applications (AED50).

Keywords: Traffic flow modeling · Traffic operations · Control · Automated vehicles · Artificial intelligence

1 Introduction

Artificial Intelligence (AI) models are being utilized extensively in different scientific and engineering domains for both analysis and predictive purposes. In particular, AI models are leveraged for processing sensor data, controlling automated vehicles (AV) and operating traffic control devices. However, there are still many challenges in those AI applications, including how to choose, build, and train AI models to avoid issues such as overfitting; translate AI models trained on synthetic (e.g., simulated) data to real-world applications; and teach AI controlled AVs how to collaborate (instead of solely maximizing their own benefits) with each other, human-driven vehicles, and traffic control devices at both local and network levels so that the overall transportation system's safety and mobility are maximized.

Among the scientific and engineering domains using AI models, automotive makers are adopting AI techniques in order to automate the movement of driver-less cars thus creating safer and more reliable automated vehicles (AV). On the other hand, traffic engineers are adopting AI to predict congestion and collision formation on our roadway networks offering real-time information for users to make better travel decisions. However, simply adopting AI instead of standard traffic flow models may lead to the lack of understanding of physical processes and dynamics leading to poor roadway performance and may produce false predictions in traffic states especially given the need of extensive data for AI training, calibration and validation purposes. Accordingly, the suitable use of AI in traffic operations and AV models requires studying two dimensions: 1) the type of AI models being adopted and their corresponding characteristics; and 2) the gap between the data available for transportation professionals and the data needed to train AI models.

Towards studying the AI model characteristics and the corresponding data needs, the Transportation Research Board (TRB) Committee on Traffic Flow Theory and Characteristics (ACP50) and the TRB Committee on Artificial Intelligence and Advanced Computing Applications (AED50) organized a breakout session at the Automated Vehicles Symposium 2020 (AVS20) - held virtually on July 27–30, 2020. The breakout session titled “Artificial Intelligence for Automated Vehicle Control and Traffic Operations: Challenges and Opportunities” brought together six scholars from academia and the industry. These scholars presented their latest work in AI as related to the traffic engineering and AV field. Following the presentations, a panel consisting of five of the invited speakers had extensive discussion with the audience. This chapter summarizes these presentations while identifying the key challenges in adopting AI for traffic and AV modeling and the corresponding efforts made to adapt data for training and calibration purposes. In particular, the objectives of the session are to:

- Identify the opportunities and challenges associated with AI applications in AV control and traffic operations

- Propose solutions for addressing the challenges
- Identify innovative applications made possible by AI and AV
- Explore how AI can enable collaborative behaviors and their impacts on transportation

Towards realizing the aforementioned objectives, the remaining sections of this chapter are organized as follows: Sect. 2 presents a summary of the 6 invited talks and Sect. 3 introduces the key results from the panel discussion.

2 Research on Utilizing Artificial Intelligence (AI) for Traffic Operations and Automated Vehicle (AV) Modeling

This section presents a summary of the six invited talks, which addressed the research challenges, opportunities and existing efforts in adapting AI to design better Automated Vehicles (AV) and to capture their impact on traffic flow. The summary includes the motivation and contributions associated with the presented research, the main conclusions, and future research directions.

2.1 The Value of Good Old-Fashioned Parametric Models for AV Control¹

Adopting artificial intelligence (AI) and self-learning algorithms (SLAs) have had a significant modelling impact on the development of vehicle automation systems. Despite having great potential, such adoption has some limitations. AI and SLAs are usually data hungry and do not behave well in new situations without proper training. This can lead to major safety issues when deploying Automated Vehicles (AV). As long as no extensive data repositories are provided to AV developers covering a wide range of traffic conditions, there will be misalignments between vehicle dynamics/movements controlled by SLAs and human driving behavioural adaptations. Such misalignments might be caused by (1) AI anomalies leading to unpredictable “harmful” movements; and (2) interactions between human and AV systems/interfaces due to lack of training and communication. As parametric models tend to capture the underlying human driving behaviour with specific modelling and theoretical constructs, they can play a role to mitigate such misalignments. A three-pronged approach (Melman et al. 2020) might be suggested to address this problem: *i*) mitigating misalignments by modelling realistic driving behaviour; *ii*) including parameterized models of driver behaviour adaptation into interaction design; and *iii*) offering human-centred interaction design.

i) Modelling realistic driving behaviour: Driver behavioural understanding is essential to modelling realistic driving behaviour. Recently a quantification of Gibson’s safe field of travel has been proposed (Kolekar et al. 2020a) as the underlying principle for a generalizable driver model. In order to compute a perceived risk, this theory evaluates the consequences/utilities of events occurring in the driving scene and the driver’s subjective belief related to the probability of an event to occur. Combined with an assigned weight, the model quantifies the perceived risk and is able to describe and predict different naturalistic driving behaviours in various traffic scenarios (Kolekar et al., 2020b).

¹ By David A. Abbink, Delft University of Technology, Netherlands.

One major benefit of this model is that it can perform well in unobserved situations with no readily available data. Accordingly, this modelling approach can contribute to mitigating one of the major problems associated with AI based autonomy.

ii) *Including models of driver behaviour adaptation into interaction design*: Drivers might adopt undesirable or risk seeking behaviour when using SAV (Semi-Automated Vehicles) or AV. Melman et al. (2017) showed in their study that drivers tend to drive faster with the haptic lane keeping assistance system. To mitigate such type of emergent behaviour, in-depth human factor and behavioural adaptation studies are necessary.

iii) *Offering human-centred interaction design*: researchers may offer breakthroughs associated with AV development and deployment; however, even when the proposed AV systems do offer a perfect safety record and a significant efficiency improvement, there will always be need for human-automation interaction. Human-automation interaction can be categorized into two categories: traded control and shared control. In traded control, at any specific time during the driving event, either the algorithm or the human controls the vehicle. This approach is comparably easy to implement and computationally less complex. In shared control, human and algorithm can both control the vehicle at any given time. Abbink et al. (2012) demonstrated one such system in which torques on the steering wheel is used for the interaction between the human driver and the algorithm. This torque is used to inform human driver about the disagreement between the trajectories produced by the algorithm and human. In the aforementioned study, the time to lane crossing (TLC) is measured for a human controlled vehicle, a shared controlled vehicle and a traded controlled (automated) vehicle through a simulator environment and it is found that the shared controlled vehicle always performs better. Moreover, in the case of automated system failure, the shared controlled system performs better because it takes less time for humans to take over the control of the vehicle and to react to the situation at hand (if compared to the traded controlled system).

In conclusion, the “old-fashioned” parametric models can play a critical role in solving misalignments between self-driving algorithms and humans, while adapting the human decision-making process to the automation technology and increasing AV safety.

2.2 Deep Learning Based Eco-driving for Connected and Automated Vehicles²

Human and freight transportation is one of the most energy-consuming sectors in the United States (US). According to a survey by the Energy Information Administration (EIA), about 28% of the total US energy consumption was associated with the transportation sector in 2019. To develop a more energy efficient and sustainable transportation system, the Connected and Automated Vehicle (CAV) technology emerges as one of the transformative solution approaches to such an environmental problem. Connected eco-driving refers to the connected and automated driving system that seeks to minimize the expected total vehicle energy consumption by taking optimal and valid actions. This system also takes into consideration other factors such as maximum travel time, fuel and battery cost. Current research in energy efficient vehicles modelling can be divided into three categories based on the methods used: rule-based models, optimization-based models, and deep-learning models. Usually, the rule-based models are simple to implement,

² By Peng Hao, University of California Riverside, U.S.A.

computationally efficient and suitable for real-time use. However, they are designed based on the assumptions and the experience of the researchers and cannot guarantee that the solution is the best (optimal) strategy among all possible alternative strategies. The optimization-based models can define objectives and search and find the best (local or global) alternative strategy. However, these models are computationally complex and most of them are not suitable for real-time and real-world implementation. Moreover, optimization-based models often underestimate the impact of exogenous aspects of the driving environment (i.e., not considered in the formulation either in the objective function or in the constraints).

To address the limitations of the optimization-based models, a graph based modular hybrid model is introduced by Hao et al. (2015) which uses graph theory models and learning-based modules. This model adopts a different approach (including machine learning algorithms) for each of the modules to archive the optimal speed and trajectory plan for energy efficient driving. It includes long short-term memory (LSTM) based signal timing prediction, radial basis function neural-network-based speed forecasting, machine learning based trajectory planning algorithms (MLTPA), etc. for real-time and effective execution of the model.

Deep learning-based modules utilize different deep learning algorithms for connected eco-driving. There are three different logical tasks associated with such modules: *i*) energy efficiency, *ii*) interaction with other traffic units, and *iii*) interaction with infrastructure. One of the challenges encountered is the implementation of these three different logical tasks in the same deep-learning construct. Hao et al. (2020) introduce a hybrid reinforced learning-based approach for eco-driving at a signalized intersection. Markov Decision Process (MDP) is used in their study to address the challenge of implementing three logical tasks in one problem. MDP is a mathematical framework that can be used to model decision making based on the interaction between the learning agent and its environment.

Dueling Deep Q Network (DDQN) is found to be the best among all neural networks studied by Hao et al. (2020). The agent vehicle has on-board sensors for knowing its current state as well as the surrounding traffic environment. It receives V2I (vehicle to infrastructure) information using Dedicated Short-Range Communication (DSRC) system or 5G cellular data. An on-board computer equipped with the decision manager algorithm calculates the long short-term reward of an action (to maximize an objective function over the whole trip instead of the immediate next few steps). The neural network model proposed by Hao et al. (2020) has two main components: a hidden feature extraction component and a policy network which is based on the DDQN architecture. Unity 3D is used to create a virtual reality environment for testing the proposed system. Three types of vehicle agents (governed by three models) are implemented in the virtual reality simulation: an intelligent driver model vehicle, a fast-speed model vehicle (always seeking to maximize the speed) and the eco-driving model vehicle based on DDQN. Results from the study show that the DDQN deep learning model vehicle performs better in terms of energy efficiency. The vehicle also has a smoother acceleration and deceleration pattern and an improved lane-changing performance if compared to the other two vehicles.

In conclusion, the deep learning-based model shows great potential in developing eco-driving strategies for CAV. CAV in eco-driving mode can significantly reduce energy consumption in the transportation sector and help move towards a more sustainable transportation system.

2.3 Machine Learning Methods – Beware: There is Nothing to Learn About Congested Urban Networks³

Congested urban networks have long been considered to behave chaotically and to be very unpredictable. This apparent complexity has led to the development of numerous signal control algorithms, mathematical programs and learning-based control methods to optimize network performance. Most of the research shows some operational improvements but they mostly correspond to light traffic conditions or very specific small networks. The recent empirical verification of the existence of a network-level Macroscopic Fundamental Diagram (MFD) suggests a different result when studying congested networks. The network MFD is a way of describing the traffic flow of urban networks at an aggregate level which is used for displaying network simulation output in a concise way. Though the turning probability at intersections is also a key variable that significantly affects the MFD, it is not well understood in the research. Moreover, there is a gap in the deep reinforcement learning (DRL) literature associated with the analysis of the different aspects of large traffic flow networks that influence the performance of DRL methods. It is not clear if and how network congestion levels affect the learning process, nor if other machine learning methods are effective, nor if current findings also apply to large networks.

Laval and Zhou (2020) provide additional evidence for the congested network property of the MFD and analyze how these properties affect the performance of machine learning methods applied to signal control. The traffic flow model used in this study is a cellular automaton (CA) implementation of the kinematic wave model with a triangular flow-density fundamental diagram, which is the simplest model able to predict the main features of traffic flow. A grid network of bidirectional streets with one lane per direction and with a traffic light at all intersections is used as the simulation network. To attain spatial homogeneity, the network is defined on a torus where each street can be thought of as a ring road where all intersections have 4 incoming and 4 outgoing approaches. Vehicle routing is set to random. A driver reaching the stop line, will choose to turn with probability p or keep going straight with probability $1 - p$. Traffic signals in the simulation operate under the simplest possible setting with only red and green phases (no lost time, all-red, yellow nor turning phases). All the control policies considered are incremental in the sense that decisions are taken every g time steps, which can be interpreted as a minimum green time: After the completion of each green time of length g , the controller decides whether to prolong the current phase or to switch light colors.

The baseline experiment shows that urban networks are more predictable than previously thought with respect to signal control and the network throughput is independent of traffic signal control even for inhomogeneous networks. To analyze the performance of AI methods, three methods are used for training the signal control policy: random

³ By Jorge Laval, Georgia Institute of Technology, U.S.A.

search, supervised learning and DRL. The random search method shows that all policies, no matter how inefficient, are optimal when the density exceeds approximately 75% indicating that the network throughput is independent of traffic signal control. Supervised learning training the policy with only two examples yields a near-optimal policy. The simulations indicate that DRL policies are only competitive and lose their ability to learn a sensible policy as the training density increases. Such a finding also indicates that the more the congestion is, the less the policy affects intersection throughput.

The main takeaway from the study is that, on congested urban networks, intersection throughput tends to be independent of signal control. It can be conjectured that this prevents DRL methods from finding sensible policies under congested conditions. In other words, all the DRL methods proposed in the literature to date may be unable to learn sensible policies and may deteriorate as soon as congestion appears on the network.

2.4 On the Challenges of Building a Camera-Only, Complete, Self-driving System⁴

Current technologies have either sophisticated technologies with low accuracy requirement or simple technologies with high accuracy requirements. Automated Vehicle (AV) is both a sophisticated technology and requires extremely high accuracy detection. For human drivers, the accidents that involve injuries and/or fatalities occur approximately every 10^4 h and 10^6 h respectively. Accordingly, to improve safety, AV should have a mean time between failures (MTBF) to be at least 10^7 h. The challenge is then to achieve such a high accuracy AV system and to validate such a system with appropriate data. The AV system has three phases: sensing, planning and acting. In the sensing phase, the AV system, with the help of different sensors, builds a three-dimension (3D) environment surrounding the corresponding vehicular space. In the planning phase, it analyzes such environment and finds the optimal driving strategy. In the acting phase, it executes the actions identified in the planning phase.

AVs need very high accuracy sensing for the required large MTBF (needed for safety considerations). In order to tackle such sensing challenge, the concept of redundancy is used. The approach is to build two fully independent subsystems, one with only cameras and another with radars and lidar. These two independent sub-systems should aim to reach a MTBF of 10^4 h each. If the two sub-systems are truly independent, even in the worst-case scenario (with a MTBF of $10^{3.5}$ h), the sensing technology can achieve the safety standards of 10^7 h MTBF. The objective becomes building an only camera-based subsystem that can reach a MTBF of 10^4 h. The main challenges in camera-based sensing come from the fact that cameras are inherently two-dimension (2D) systems and yet we need a 3D understanding of the surrounding environment with high accuracy even in edge cases (where visibility is very low). There are several methods to produce 3D data from 2D camera like prediction of object dimension, visual lidar (VIDAR)/structure from Motion, etc. VIDAR uses deep learning to generate a 3D model of the environment from the camera feed. Another approach is to project the 3D map information into a 2D image plane and then use the 2D data from the camera for the planning phase. The

⁴ By Shai Shalev-Shwartz, Mobileye.

redundancy approach uses multiple detection methods like VIDAR, scene segmentation 3DVD, etc. and multiple measurement techniques.

The planning phase involves decision making to find optimum actions to avoid accidents. However, to find actions avoiding accidents at all cost may not be an ideal solution for this phase. There should be a balance between being a perfect driver and being a driver who blends in. This also raises some ethical questions. The usefulness/safety tradeoff adopted by humans requires a sense of caution. The duty of care (Tort Law) states that a legal obligation is imposed to an individual rewiring adherence to a standard of reasonable care while performing any act that could foreseeably harm others. Human has common sense to interpret such standard/law. The challenge is to interpret the law for the AV systems. Rigorous mathematical modeling is required to formalize an interpretation of any law which is applicable to AVs. The resulting models should be sound, useful and applicable to machines. Soundness implies that the interpretation by the model should comply with the common sense of human driving. The interpretation should lead to zero accidents in a utopic world. Usefulness ensures that AVs don't block traffic being over-cautious and non-agile. Efficiency of AV models should be verified for machine applicability. This is not trivial due to potential of butterfly effect. Mobileye proposes a system called Responsibility Sensitive Safety (RSS) model, which is a mathematical model to formalize a commonsense interpretation of the duty of care. RSS should provide mathematical guarantees for AV to never cause an accident, to be relevant to human drivers and to be efficiently verifiable.

In summary, to tackle the challenges of achieving AV safety standards, Mobileye focuses on 1) redundancy for 3D sensing, and 2) formal safety modeling during planning (i.e., planning phase) while considering human judgement considerations.

2.5 Mixed Autonomy Traffic: A Reinforcement Learning Perspective⁵

We imagine that the future of the transportation sector relies on fully automated and highly efficient transportation systems. It is predicted that by the year 2050 we will achieve full autonomy for surface vehicles. We have multiple billion-dollar corporations racing for the creation of the first fully automated vehicle (AV) and they are improving year after year towards reaching such a goal. There are many tools available for analyzing a single AV with full autonomy while adopting deterministic models with no uncertainty/error. The operation of a single AV depends however on other vehicles in the system and there is a need for additional studies on its impact on the whole mixed (i.e., with the existence of both automated and human driven vehicles) system. In other words, there exists a significant challenge represented by the understanding of and modeling the mixed autonomy state of the transportation system. This challenge is due to the existence of many sources of uncertainty including partial observation, limited communication, data collection challenges, etc.

Mixed autonomy can take different forms like advanced driver assistance systems. The impact of such mixed autonomy on safety, reliability efficiency, fairness should be analyzed. Understanding the impact of mixed autonomy on broader societal system is also necessary. All of these issues require more analysis tools. To analyze the problem

⁵ By Cathy Wu, Massachusetts Institute of Technology, U.S.A.

at hand, deep reinforcement learning (DRL) is used. In this modeling approach, agents are the vehicles that are automated and everything else is considered as the exogenous environment. The agents will make decisions such as when to accelerate or decelerate according to a learned policy in order to maximize reward. The average velocity, energy consumption, travel time, safety and comfort should be considered in the global reward function. Ultimately the goal is to study large urban networks where a fraction of the vehicles is automated. Wu et al. (2017) explores the potential of DRL methods when training the algorithm from scratch. This study designs a representative set of scenarios that exhibits a variety of different traffic phenomena including intersection, bottleneck and on/off ramp scenarios. By using DRL with 5–10% of AVs, the simulations show from 30% to 142% increase in average velocity across the scenarios. Some of the learned policies match the performance and behavior of the control strategies devised by experts over the years. These findings provide validation of the methodology to analyze the impact of mixed autonomy in urban environments. A critical challenge of analyzing these systems in large-scale contexts comes from the fact that no two cities' traffic networks are the same; even the traffic network in a single city varies from block to block. There is a combinatorial number of environments that exists when it comes to traffic networks. Accordingly, the approach of training for each scenario will not be practical moving forward. A potential solution to this challenge is to use transfer learning. Transfer learning is the use of knowledge gained from a source task to bias the learning process on a target task while forming a set of good hypotheses. A zero-shot transfer is where no learning is done on the target task and is analogous to out-of-distribution generalization in supervised learning. Kreidieh et al. (2018) investigates the transferability of knowledge from a circular source environment to an open street network environment and shows that knowledge transfer is possible between these two sources using zero-shot transfer. Ongoing research is looking into the possibility of learning from a single policy and applying the findings on many different scenarios.

Ultimately, the goal is to develop a set of techniques that can help analyze the mixed autonomy in the existing surface transportation system. This requires rigorous studying of mixed autonomy systems; DRL is a promising technique for the resulting modeling toolkit. However, there is a long way to go to build a toolkit for analyzing the whole urban system and translating the modeling results based on simulation to support real-world urban planning decision in different network architectures.

2.6 Traffic State Estimation with Physics Regularized Machine Learning: A New Insight into Machine Learning Applications in Traffic Flow Modeling⁶

Traffic state estimation (TSE) is the precursor of a variety of advanced traffic operational tasks. As the traffic sensors on freeway networks can only cover a limited range of areas, TSE is a useful tool to provide full-field traffic information. TSE models estimate flow rates and speeds over the whole network. Most TSE models in the literature are derived from macroscopic traffic flow models. In the early research stages, macroscopic traffic dynamics are found to be similar to hydrodynamics. The associated models are formulated for ideal conditions and significant effort is needed for their calibrations.

⁶ By Terry Yang, University of Utah, U.S.A.

It is also difficult to work with the noisy and fluctuated data collected by traffic sensors. To deal with the noise and fluctuations, stochastic traffic flow models are used. These models may be divided into two types: stochastic extensions and stochastic formulations. Stochastic extension models add Gaussian noise to the model expression in order to quantify the noise from the sensor data. However, they can produce mean dynamics that do not coincide with the original deterministic dynamics due to non-linearity. Stochastic formulation models do not have this inconsistency problem; however, they might lose the ability to obtain a mathematical solution due to the lack of a closed form expression/methodology. With the increase in data availability, researchers in recent years started to look into more data-driven machine learning (ML) approaches. In general, the data-driven ML models can outperform the classical traffic flow models; however, the performance of these models still heavily relies on the quality and the quantity of available data. In order to mitigate such limitation, Physics Regularized Machine Learning (PRML) is introduced. PRLM is a novel modelling framework which encodes the classical traffic models (“physics models”) into the ML framework: output from the physics models is used to later train the regularized ML model to improve the model performance. If compared to classical traffic flow models, PRML can effectively capture data uncertainty and reduce the efforts associated with model calibration. If compared to ML models, PRML is more robust as it better handles noisy training data (through a Gaussian process – GP) and is more explainable in terms of model performance.

Yuan et al. (2020) develop a stochastic physics regularized Gaussian process (PRGP) which uses a Bayesian inference algorithm to estimate the mean and the kernel of the PRGP. A physical “regulator” based on macroscopic traffic flow models is also developed to augment the estimates via a shadow GP and an enhanced latent force model is used to encode physical knowledge into stochastic processes. Based on the posterior regularization inference framework, an efficient stochastic optimization algorithm is developed to maximize the evidence lower bound of the system likelihood. The model is evaluated using four detector data sets from I-15 in Utah, US. The results from the case study show that all four PRGP models perform better than some standard ML models in terms of flow estimation and produce comparable and acceptable results in terms of speed and travel time estimation. Moreover, the findings show that encoding a bad physics model into the PRGP can downgrade the model performance. PRGP models also produce better estimates than the physics models. To study the robustness of the proposed PRGP model, artificial noise is added to the training dataset while keeping the test dataset constant. With noisy data, PRGP models produce acceptable estimates of flow, speed and travel time and the errors from this type models are less than the other ML and physics models. Adding more sensor data into the training dataset further improves the models’ performance.

The PRGP can greatly outperform physics models in capturing the data uncertainty and fluctuations. When the training dataset is sufficient and accurate, PRML only slightly outperforms standard ML models in terms of speed, flow and travel time estimation. With noisy data, PRML is more robust than the ML and physics models. It can be noted that encoding a more advanced physics model can help the PRGP produce better estimates, while encoding an inapplicable physics model can downgrade the model performance.

3 Discussion

Advancements in the sensing and computational technologies have made the collection and the utilization of big data to solve real-world transportation problems more practical. In recent years, numerous research efforts have been seen in utilizing data-driven and AI methods for automated vehicles modelling and traffic operations. Such research efforts have led to the following question: what will be the role of core traffic scientists in the transportation sector and are traditional traffic flow and traffic operations models still relevant? Though applying AI approaches doesn't need the fundamental understanding of traffic science, traffic researchers need to well define the problems at hand. Defining the problems' characteristics and boundaries is a key to develop efficient and reliable AI models to solve traffic problems. The fundamental understanding of the mixed traffic flow dynamics and network infrastructure and control specifications is needed to adopt AI models especially when modeling AV systems and estimating traffic states. Complex neural network models are not always required and do not always perform better. In fact, in some cases (e.g., when developing network level optimum signal control strategy) simple supervised machine learning algorithms perform better.

The ultimate goal for the future transportation system is to make the whole transportation system fully automated with an increased level of safety, reliability, efficiency, and sustainability. Towards achieving such a goal, present researchers focus on individual autonomy, connected and automated systems and their network level performance. In line with such research directions, this breakout session presents six research efforts.

The keys findings from the presentations and the subsequent discussion were:

- Human interaction with AVs might play a significant role in AV system performance in the future and it should be incorporated in the AV system design.
- The transportation sector is one of the major energy consuming sectors in the US and connected AV systems can play a significant role to increase the energy efficiency and thus decrease energy consumption.
- Supervised learning can have great potential for AV control and traffic operations. Particularly, the traffic domain expertise can be used to design the problems faced, choose the proper AI or ML techniques to be adopted, evaluate the performance of the proposed methods, and interpret the corresponding results.
- Safety has been the upmost important factor in AV regulations and product developments. However, there needs to be a balance between safety and other performance aspects (e.g., congestion, energy consumption, fairness). Currently, safety remains the primary concern of the AV industry and the focus may be expanded to address other metrics/dimensions in the next 5 to 10 years.
- There is a lack of available toolkits to analyze the impact of mixed autonomy. Understanding mixed autonomy can help make better policy to smoothen the transition process from manually operated traffic components to fully automated transportation system.

The aforementioned findings motivate the following research needs/outcomes and may guide future research associated with adopting AI in AV modeling/development and traffic operations:

- It is important to devise regulations/policies to guide the development and the adoption of AV and AI technologies (e.g., those related to safety and ethics). This is a challenging but a very important task with direct impact on the humanity.
- Two questions were raised regarding the role of data in adopting AI: (1) how much data is needed to evaluate the performance of a model? (2) How much data is needed/enough for AI training? The first question has been well studied. Although it is generally agreed that more data is beneficial for AI models' training/learning, the "enough" part of the second question has not been well addressed. Overall, more data collection and sharing efforts are needed.
- More collaborations are needed among government, academia, the AV industry, etc. For example, if researchers do not understand how AVs function, they will find difficulties in thoroughly evaluating the AVs' impacts on surface roadway network performance. Similarly, AI experts and transportation engineers should work closely to better address practical problems encountered on a daily basis in complex environments.
- Experts from academia and the industry need to connect with the policy makers to make informed decisions.
- Joint research between the government and the AV industry is needed to develop standards associated with insurance, security and communication strategies (e.g., vehicle to infrastructure -V2I- and vehicle to vehicle -V2V- communication standards).

Acknowledgements. The authors would like to acknowledge the remaining breakout session organizers: Alireza Talebpour (University of Illinois at Urbana Champaign), Mecit Cetin (Professor, Old Dominion University), Mehdi Hashemipour (Data Scientist, US Department of Transportation), Mo Zhao (Research Scientist, Virginia Transportation Research Council), Mohamed H. Zaki (Assistant Professor, University of Central Florida), Pan Lu (Associate Professor, North Dakota State University), Simeon Calvert (Delft University of Technology), Steven Mattingly (Professor, University of Texas at Arlington), and Xiaoyu (Sky) Guo (PhD Student, Texas A&M University – College Station). Additional thanks go to the TRB Committee Section/Committee chairs for their leadership and sponsorship of the session: Robert Bertini (Oregon State University), Soyoung Ahn (University of Wisconsin, Madison), and Sherif Ishak (Old Dominion University). This book was made possible with extensive outreach and synthesis efforts while coordinating the event details with the AVS2020 organizing committee.

References

- Abbink, D.A., Mulder, M., Boer, E.R.: Haptic shared control: smoothly shifting control authority? *Cogn. Technol. Work* **14**(1), 19–28 (2012)
- Hao, P., Wu, G., Boriboonsomsin, K., Barth, M.J.: Developing a framework of eco-approach and departure application for actuated signal control. In: 2015 IEEE Intelligent Vehicles Symposium (IV), pp. 796–801. IEEE (2015)
- Hao, P., Wei, Z., Bai, Z., Barth, M.J.: Developing an adaptive strategy for connected eco-driving under uncertain traffic and signal conditions. UC Davis: National Center for Sustainable Transportation (2020). <https://doi.org/10.7922/G2F18WZ1>
- Kolekar, S., de Winter, J., Abbink, D.: Which parts of the road guide obstacle avoidance? Quantifying the driver's risk field. *Appl. Ergon.* **89**, 103196 (2020a)

- Kolekar, S., de Winter, J., Abbink, D.: Human-like driving behaviour emerges from a risk-based driver model. *Nat. Commun.* **11**(1), 1–13 (2020b)
- Kreidieh, A.R., Wu, C., Bayen, A.M.: Dissipating stop-and-go waves in closed and open networks via deep reinforcement learning. In: 2018 21st International Conference on Intelligent Transportation Systems (ITSC), pp. 1475–1480. IEEE (2018)
- Laval, J., Zhou, H.: Congested urban networks tend to be insensitive to signal settings: implications for learning-based control. arXiv preprint [arXiv:2008.10989](https://arxiv.org/abs/2008.10989) (2020)
- Melman, T., De Winter, J.C.F., Abbink, D.A.: Does haptic steering guidance instigate speeding? A driving simulator study into causes and remedies. *Accid. Anal. Prev.* **98**, 372–387 (2017)
- Melman, T., Beckers, N., Abbink, D.: Mitigating undesirable emergent behavior arising between driver and semi-automated vehicle. arXiv preprint [arXiv:2006.16572](https://arxiv.org/abs/2006.16572) (2020)
- Wu, C., Kreidieh, A., Parvate, K., Vinitsky, E., Bayen, A.M.: Flow: Architecture and benchmarking for reinforcement learning in traffic control. arXiv preprint [arXiv:1710.05465](https://arxiv.org/abs/1710.05465), October 2017
- Yuan, Y., Yang, X.T., Zhang, Z., Zhe, S.: Macroscopic traffic flow modeling with physics regularized Gaussian process: a new insight into machine learning applications. arXiv preprint [arXiv:2002.02374](https://arxiv.org/abs/2002.02374) (2020)