Semper in Motu: Transforming Mobility Through Learning and Control

2024 BODE PRIZE LECTURE

KARL H. JOHANSSON

Dear ladies and gentlemen, colleagues, and friends. Let me start by thanking IEEE Control Systems Society (CSS) President Magnus Egerstedt for the kind introduction and thanking the members of the Bode Award Committee for this amazing honor. In addition, it is a privilege to have the opportunity to talk about some of the research that I am currently working on (see "Summary"). Congratulations also to the 2024 IEEE Conference on Decision and Control organizers; it has truly been a fantastic week here in Milan.

The Bode Lecture has accompanied me since the beginning of my career. I started my Ph.D. studies at Lund University in 1992, just a few years after Gunter Stein gave the first Bode Lecture (see "CSS Hendrik W. Bode Lecture Prize"). It was titled "Respect the Unstable" and was unusual at the time because it was recorded. My Ph.D. advisor, Prof. Karl Åström, had a VHS cassette copy of the recording. Watching this video was compulsory for Ph.D. students at Lund University, in Sweden. Despite the video quality of the recording being subpar compared to today's standards, the lecture was beautiful, and it was easy to understand why it has become a classic in the control systems community and beyond. Dr. Stein used the example of a very simple flight control system and showed that there exist fundamental limitations on what feedback can achieve in practice related to the location of the system's right half-plane poles and zeros; the lecture convincingly illustrated how control theory

KARL H. JOHANSSON GIVING THE BODE PRIZE LECTURE AT THE 2024 IEEE CONFERENCE ON DECISION AND CONTROL, MILAN CONVENTION CENTRE, ITALY.

Digital Object Identifier 10.1109/MCS.2025.3593108
Date of current version: 18 September 2025

Summary

The transformation of the transport sector is continuously in motion. Through advances in sensing, connectivity, computing, and electrification, the control community has been and will continue to be actively engaged in the shaping of a sustainable and efficient infrastructure for moving people and goods. Although self-driving technologies have garnered significant attention, achieving widespread and safe deployment remains a challenge. Meanwhile, innovations are continuing to occur to optimize and improve the resilience of transport systems, highlighting the broader impact of control technology on mobility. This lecture explores intelligent transport and the influences of the emerging field of cyberphysical–human systems, using three case studies: 1) vehicle automation and occlusion, 2) traffic control using physics-informed machine learning, and 3) rollout-based planning for electric truck charging coordination. These case studies are used to illustrate how uncertainty can be represented and mitigated using safety-first reachable set computations, a mix of physics- and data-driven models, and stochastic forecasting with dynamic programming (DP). The article highlights joint work with students, postdoctoral researchers, and collaborators in academia and industry.

(in this case, Bode's sensitivity integral) has had a significant impact on the understanding of what engineering can do (how to build flight control systems).

I am happy to recommend that all the students in the audience watch this video. My advisors, Profs. Åström and Rantzer, emphasized in the very same spirit to me that the development of good theory should be motivated by applications, and conversely, good methods being developed should have an impact on real systems. This is something I have carried with me throughout the research we have been conducting in my group. At the core of our approach is the mitigation of uncertainty through feedback control at multiple layers of transportation systems. This is the main message that I would like to convey in this article.

CSS Hendrik W. Bode Lecture Prize

The CSS Bode Lecture Prize recognizes distinguished contributions to the field of control systems science or engineering. The basis for judging is the technical merit of the contribution to control systems science or engineering and also includes the broader impacts of the contribution toward the benefit of society at large and other relevant aspects. The recipients, with their quotations and photos, appear on the CSS home page, https://ieeecss.org/awards/ieee-css-hendrik-w-bode-lecture-prize, and are as follows:

- 1989: Gunter Stein
- 1990: David Luenberger
- 1991: Petar Kokotovic
- 1992: Brian Anderson
- 1993: Michael Athans
- 1994: Gene Franklin
- 1995: Bob Narendra
- 1996: Jürgen Ackermann
- 1997: Edward Davison
- 1998: J. Boyd Pearson
- 1999: Graham Goodwin
- 2000: Mathukumalli Vidyasagar
- 2001: Alberto Isidori
- 2002: Eduardo Sontag
- 2003: Lennart Ljung
- 2004: Tamer Başar
- 2005: Pravin Varaiya
- 2006: Arthur Krener
- 2007: P.S. Krishnaprasad
- 2008: Christopher Byrnes
- 2009: Peter Caines
- 2010: Manfred Morari
- 2011: John Baillieul
- 2012: Jessy Grizzle
- 2013: B. Ross Barmish
- 2014: Bruce Francis
- 2015: Hassan Khalil
- 2016: Richard M. Murray
- 2017: Naomi Leonard
- 2018: Mark Spong
- 2019: Lei Guo
- 2020: Kristin Y. Pettersen
- 2021: Pramod Khargonekar
- 2022: David J. Hill
- 2023: Miroslav Krstic
- 2024: Karl H. Johansson

THE TRANSFORMATION OF THE TRANSPORT SECTOR

Let us first discuss why there is such a large and urgent need for a technical and digital transformation of the transport sector. The aim is to achieve zero emissions from transportation by 2070, limiting the global temperature rise. In Figure 1, we see the carbon dioxide emissions from transport in the "Sustainable Development Scenario" of the International Energy Agency. Thanks to future technical and other advances, the emissions are expected to decrease considerably. Note that the relative impact of emissions from various transport modalities is predicted to change; for example, the transportation of goods through heavy-duty vehicles may play an increasingly important role in the future.

How do we achieve this positive development to reach zero emissions from transportation? It will be an essential task for many of the students in the audience to dedicate their careers to innovate in this space. Only with new groundbreaking research, where I believe control engineering will play a crucial role, will we see this evolution

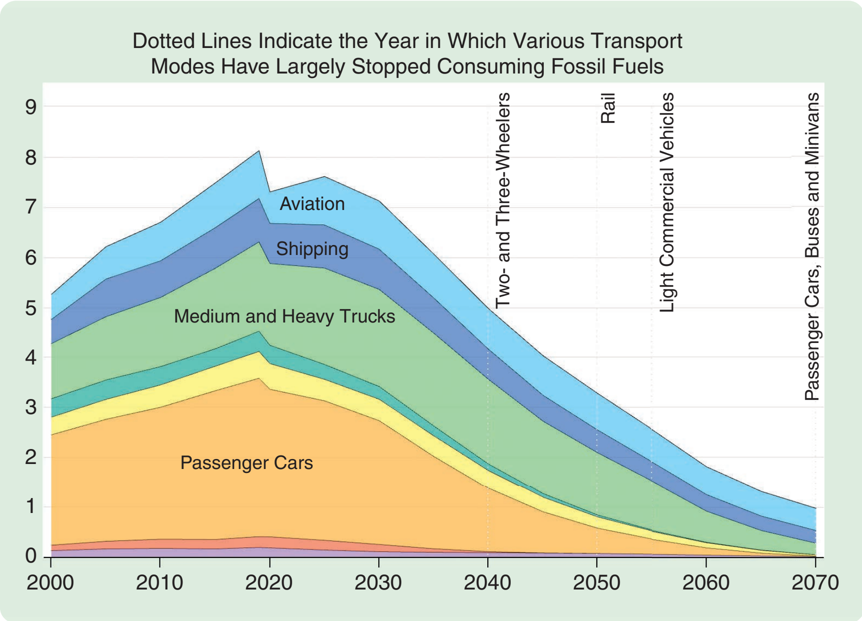

FIGURE 1. Gigatons of carbon dioxide released into the atmosphere each year for various modes of transportation from 2000 to 2070. The prediction into the future corresponds to a sustainable development scenario in which significant future technical and other advances help to reduce the emissions considerably. (Source: International Energy Agency [1]; used with permission.)

(see "*IEEE Control Systems* on Automated Vehicles and Transportation"). There are also significant commercial opportunities in this development, which the automotive industry has already recognized. Many of its R&D strategies revolve around the key areas indicated in Figure 2.

Sweden is uniquely positioned in this field, as it has a long tradition in the automotive industry, and some of the key innovations in this area were developed by Swedish engineers, such as the three-point seatbelt, which dates back to the 1950s. Volvo Group and Scania are two leading truck manufacturers internationally, both of which invest heavily in supporting long-term research collaborations with us and other research groups as well as developing experimental vehicle concepts (Figure 3).

IEEE Control Systems on Automated Vehicles and Transportation

Over the last 20 years, there have been many articles on vehicle automation and road transport in *IEEE Control Systems* magazine as well as special issues (Figure S1). The optimization of hybrid electric vehicles is discussed in [S1]. The authors of [S2] present novel truck platoon controllers and their experimental evaluations. Formal methods for controlling traffic flow are described in [S3]. Experimental testbeds with real or miniature robotic vehicles, such as the one discussed in "Small Vehicles for Autonomy: A Rapid Experimental Testbed for CAVs," are important for research and education. Another such testbed is described in [S4]. Earlier this year, there was a whole issue of the magazine dedicated to mixed-autonomy traffic [S5]. One of the articles presents an open-road field experiment with 100 connected and automated vehicles (CAVs) [S6].

FIGURE S1 Issues of *IEEE Control Systems* dedicated to vehicle automation, mixed-autonomy traffic, and related control applications.

REFERENCES

[S1] B. Egardt, N. Murgovski, M. Pourabdollah, and L. M. Johannesson Mardh, "Electromobility studies based on convex optimization: Design and control issues regarding vehicle electrification," *IEEE Control Syst. Mag.*, vol. 34, no. 2, pp. 32–49, Apr. 2014, doi: 10.1109/MCS.2013.2295709.

[S2] A. Alam, B. Besselink, V. Turri, J. MåRtensson, and K. H. Johansson, "Heavy-duty vehicle platooning for sustainable freight transportation: A cooperative method to enhance safety and efficiency," *IEEE Control Syst. Mag.*, vol. 35, no. 6, pp. 34–56, Dec. 2015, doi: 10.1109/MCS.2015.2471046.

[S3] S. Coogan, M. Arcak, and C. Belta, "Formal methods for control of traffic flow: Automated control synthesis from finite-state transition models," *IEEE Control Syst. Mag.*, vol. 37, no. 2, pp. 109–128, Apr. 2017, doi: 10.1109/MCS.2016.2643259.

[S4] B. Chalaki, L. E. Beaver, A. M. I. Mahbub, H. Bang, and A. A. Malikopoulos, "A research and educational robotic testbed for real-time control of emerging mobility systems: From theory to scaled experiments [Applications of Control]," *IEEE Control Syst. Mag.*, vol. 42, no. 6, pp. 20–34, Dec. 2022, doi: 10.1109/MCS.2022.3209056.

[S5] A. Annaswamy, "Mixed autonomy at scale [About this Issue]," *IEEE Control Syst.*, vol. 45, no. 1, pp. 5–9, Feb. 2025, doi: 10.1109/MCS.2024.3500393.

[S6] J. W. Lee et al., "Traffic control via connected and automated vehicles (CAVs): An open-road field experiment with 100 CAVs," *IEEE Control Syst.*, vol. 45, no. 1, pp. 28–60, Feb. 2025, doi: 10.1109/MCS.2024.3498552.

TRANSPORT AUTOMATION AND UNCERTAINTY

Ten years ago, the Society of Automotive Engineers (SAE) presented a vehicle automation taxonomy that defined six levels, ranging from no driving automation (level 0) to full autonomy (level 5) [2] (Figure 4). The SAE inspired a road map describing the path from human-driven vehicles to increasingly automated support systems. Even 10–15 years ago, the leadership of some major corporations, such as Google and Tesla, considered autonomous driving to be a solved problem. Fast-forwarding to today, the picture is somewhat different. We are still quite far from autonomous vehicles on every road; there are several fundamental challenges to be overcome. At the same time, it should be pointed out that great progress has been made, such as large-scale deployments of robot taxis underway in San Francisco, CA, USA, and other cities; automated buses demonstrated in the Stockholm, public transport system (see "Self-Driving Buses in Stockholm Pilot Study"); and autonomous vehicles for mining operations tested worldwide, including in the northern part of Sweden.

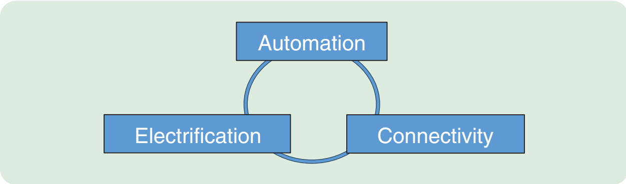

FIGURE 2. Three key areas of the automotive industry's R&D strategies.

FIGURE 3. (a) An autonomous Scania concept truck for mining applications and (b) an electric Volvo garbage truck. (Sources: Scania and Volvo Group; used with permission.)

Have we solved the vehicle automation problem? No, not really. We do have impressive experimental tests going on, but we still have a long way to go in terms of fatalities. The crash rate per million miles for automated vehicles remains more than 10 times higher than that of human-driven vehicles [3]. Healthy skepticism about building fully automated vehicles capable of driving in any road and traffic conditions was raised early by many academic leaders, including Dr. Steven Shladover, a world-leading researcher in vehicle automation and a key contributor to the creation of the intelligent transportation systems program in the United States.

Why is vehicular transportation difficult to automate? In this lecture, we argue that *uncertainty* is the key reason. Cars, buses, and trucks are driven in highly dynamic environments under drastically different conditions depending on the time of day, day of the year, location, and so on. Despite even more sophisticated sensor technologies and data collection, representation, and prediction of all possible corner cases that may arise, including cases that occur

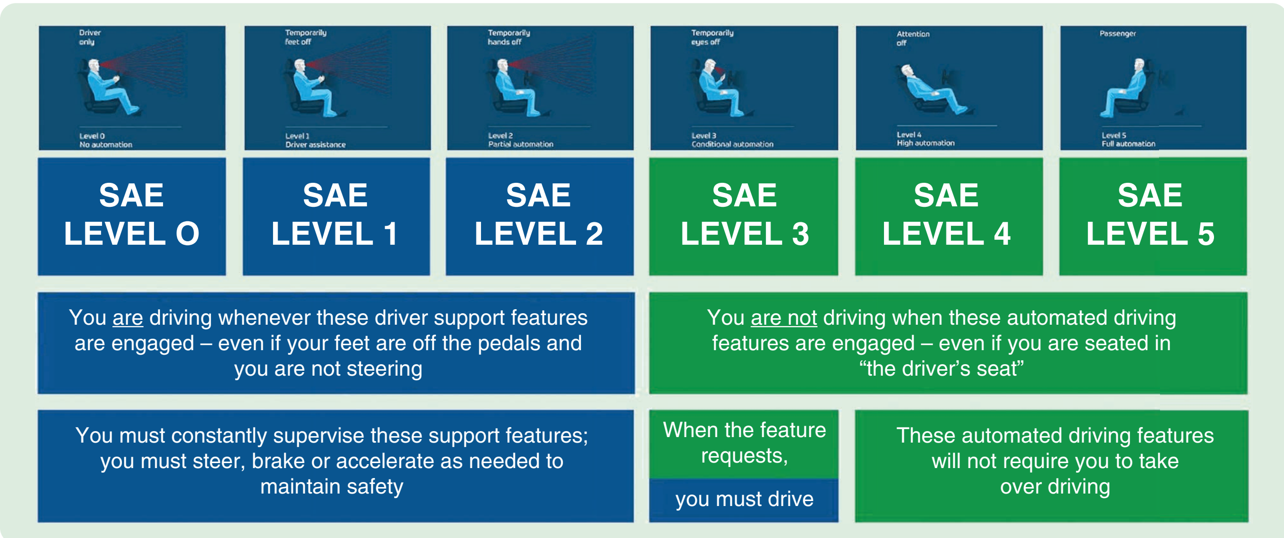

FIGURE 4. The SAE suggests six levels of driving automation, with increasing levels of automation from level 0 to level 5. The text below the levels indicates the expectation on the driver. (Source: SAE International; used with permission.)

Self-Driving Buses in Stockholm Pilot Study

Self-driving buses were tested in regular service in the Stockholm region during 2018–2023. Figure S2 shows one of the buses deployed.

Residents of the Barkarbystaden suburban area [S7] could take the self-driving bus to travel within their residential area or connect to an electric bus rapid transit (BRT) system (B), which led to a metro station (T) or a local train station (J), as indicated in Figure S3. Various technologies were evaluated in the study, including vehicle operation [S8] and human–machine interaction [S9] by KTH researchers, vehicle technology by Nobina, and 5G communication by Ericsson.

FIGURE S2 A self-driving bus integrated into the regular public transport system of the Stockholm region. (Source: Nobina; used with permission.)

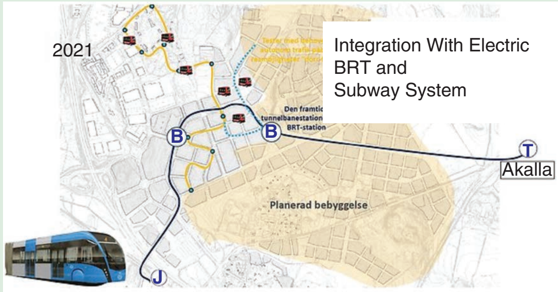

FIGURE S3 The self-driving bus picked up passengers on demand (indicated by the yellow path) and typically brought them to the BRT system (the blue path), which was connected to the wider Stockholm public transport network. (Source: Järfälla municipality; used with permission.)

REFERENCES

[S7] "Now the self-driving bus is rolling in Barkarbystaden for the last time — Almost," *Nyheter*, Sep. 29, 2023. [Online]. Available: https://tinyurl.com/yrjj2p7k

[S8] P. N. E. Chee, Y. O. Susilo, Y. D. Wong, and A. Pärneståll, "Which factors affect willingness-to-pay for automated vehicle services? Evidence from public road deployment in Stockholm, Sweden," *Eur. Transp. Res. Rev.*, vol. 12, no. 1, 2020, Art. no. 20, doi: 10.1186/s12544-020-00404-y.

[S9] A. Axelsson, B. Vaddadi, C. Bogdan, and G. Skantze, "Robots in autonomous buses: Who hosts when no human is there?" in *Proc. Companion ACM/IEEE Int. Conf. Human-Robot Interaction*, New York, NY, USA: ACM, 2024, pp. 1278–1280.

rarely, remain important and continue to be challenging. Imagine a truck running downhill in suddenly changing weather conditions and having to deal with rapidly deteriorating road surface conditions. Or consider the difference in situation awareness and driving behavior between an experienced bus driver in London, U.K., and a teenager behind the steering wheel for the very first time. Feedback control is the science of mitigating uncertainty. Hence, we believe that the control systems community has some of the fundamental tools needed to take the next major set of steps in the development of automated vehicles. This is not just about new control algorithms but also about novel architectures and providing the right information at the right time to simplify decision making (Figure 5).

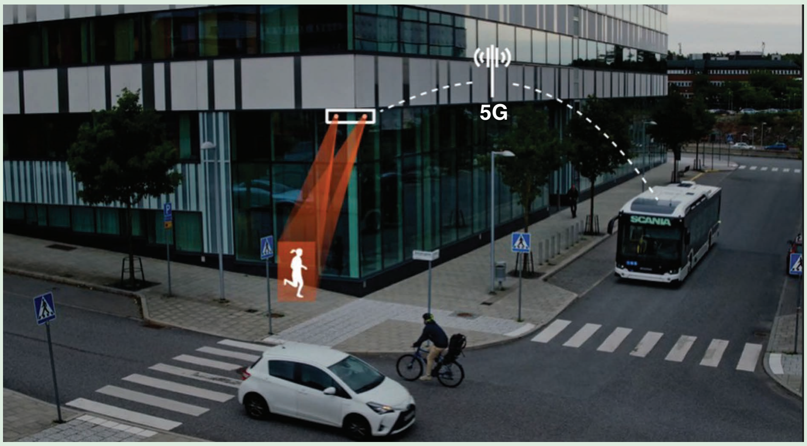

FIGURE 5. An intelligent intersection supporting safe interaction between a runner and an automated bus. A camera in a roadside unit detects the runner and communicates the information over a 5G cellular network to the bus despite the occlusion by the building, resulting in the bus slowing down and stopping at the pedestrian crossing. (Illustration from [4].)

The remainder of this lecture focuses on three case studies that illustrate aspects of how control contributes to this thesis in three "layers" of the transport system. We first demonstrate how an individual automated vehicle can handle traffic situations with occlusion. Then, we discuss how uncertain road traffic can be predicted and controlled using physics-informed machine learning models. Finally, we study rollout-based planning for electric

The development of good theory should be motivated by applications, and conversely, good methods being developed should have an impact on real systems.

truck charging coordination, considering large fleets of vehicles driving across Sweden under shared and varying resource constraints.

VEHICLE AUTOMATION AND OCCLUSION

Let us start with the first case study on vehicle automation and occlusion. Pravin Varaiya, a pioneer in control theory, gave a thought-provoking talk at a 2018 workshop in Stockholm about an Uber accident that had occurred the previous year in Tempe, AZ, USA. He described how a human-driven car turned left at an intersection while an automated Uber entered the intersection in the opposing lane. There were several other vehicles at the intersection. Neither the human driver nor the automated vehicle saw the other, resulting in a fatal collision in the middle of the intersection. The natural questions raised were how automated vehicles should handle occlusion and how they resolve the situation when vehicles have insufficient information [5], such scenarios can be tested in practice, see "Small Vehicles for Autonomy: A Rapid Experimental Testbed for CAVs").

Small Vehicles for Autonomy: A Rapid Experimental Testbed for CAVs

Connectivity has the potential to improve the safety and efficiency of automated vehicles significantly. By connecting to other vehicles, roadside units, or network infrastructure, CAVs can overcome their limitations in safety and efficiency, such as sensor occlusions or traffic congestion. However, despite the potential benefits, the validation and deployment of new CAV services have been hindered by slow and expensive testing. Many complex scenarios and rare events require systematic evaluation. One such example is illustrated in Figure S4, in which multiple CAVs and human-driven vehicles interact in a scenario where a truck has broken down and blocks part of the road. To resolve the situation, the vehicles are supported by an intelligent road infrastructure that shares information with them over multiple communication networks. A control tower with a human operator supports the vehicles in difficult traffic situations.

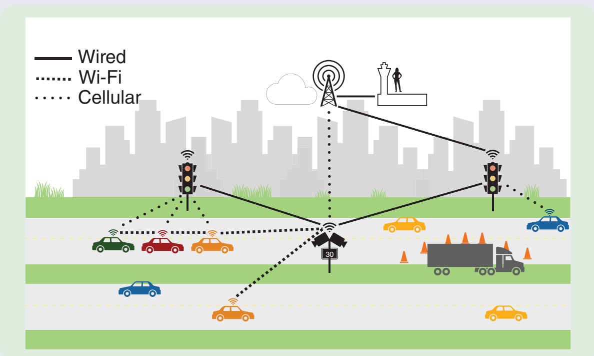

FIGURE S4 A traffic scenario with a mix of human-driven and automated vehicles supported by traffic signals, roadside units, and communication networks. Intelligence is distributed over the vehicle network to resolve most traffic situations that can appear, while a traffic control tower is prepared to assist in more complex situations.

FIGURE S5 The KTH Smart Mobility Lab SVEA platform integrated into the Kista Innovation Park, where a private 5G network is deployed to support small-scale connected vehicle experiments.

The Small Vehicles for Autonomy (SVEA) platform in Figure S5 has been developed with Swedish industry to support the rapid development and experimentation of new CAV services. The SVEA platform consists of several automated vehicles of 1/10 scale connected over a private 5G network and has been extensively used for testing and evaluation in various projects [S10], [S11].

REFERENCES

[S10] F. J. Jiang, M. Al-Janabi, T. Bolin, K. H. Johansson, and J. Mårtensson, "SVEA: An experimental testbed for evaluating V2X use-cases," in *Proc. IEEE 25th Int. Conf. Intell. Transp. Syst. (ITSC)*, 2022, pp. 3484–3489, doi: 10.1109/ITSC55140.2022.9922544.

[S11] K. M. Arfvidsson et al., "Small-scale testbed for evaluating C-V2X applications on 5G cellular networks," in *Proc. IEEE Intell. Vehicles Symp. (IV)*, 2024, pp. 149–155, doi: 10.1109/IV55156.2024.10588559.

The question of how to handle the limited information available for vehicles in certain traffic situations, which arose as a result of this fatal Uber accident, directly inspired our group's recent focus on how automated vehicles can reason with other vehicles and road users in occluded areas. Our approach [6] to model occluded areas is based on classical reachable set computations [7]. To illustrate the approach, consider an automated blue ego vehicle that is waiting at a T intersection (Figure 6). A yellow car not only hinders the ego vehicle from entering the intersection but also occludes the lane and the area behind it. A worst-case approach in this situation corresponds to the case where one or more vehicles are in the occluded area. We assume that this is the case and that any vehicle in the occluded area follows the traffic rules. It is then possible to propagate the set corresponding to the occluded area forward in time and, in this way, overapproximate the location of any potential hidden vehicle, as indicated in the figure. The propagated set can be fused with information from the ego vehicle's sensors and roadside units. This enables the ego vehicle to reason about and predict the state of traffic in the near future, thereby increasing the vehicle's so-called situational awareness, which is crucial for making informed decisions about when it is safe to drive (Figure 7).

Let us now evaluate our approach to dealing with occlusion in a slightly more complex scenario. Figure 8 demonstrates how the ego vehicle turns left at an intersection under four different levels of information and reasoning. The first two are based on onboard information, and the last two utilize information from a roadside unit mounted on a building. The dark color indicates occluded areas. The ego vehicle remains safe in all four cases, but in the fourth case, the vehicle is able to pass a longer distance due to the shared information and reasoning. For details on how to do the reachability set computations and implement the safe planner, see [6], [8], and [9]. We have also evaluated the approach to occlusion in practice at an intersection similar to the one above (Figure 9). An automated Scania truck (ego vehicle) waits at a T intersection while processing lidar scans from the ego vehicle (green) and from a roadside unit (blue). Using these scans, the truck can determine whether it is safe to drive or not. For automated vehicle experiments under more complex driving conditions, see "Automated Truck Driving Under Low-Friction Conditions").

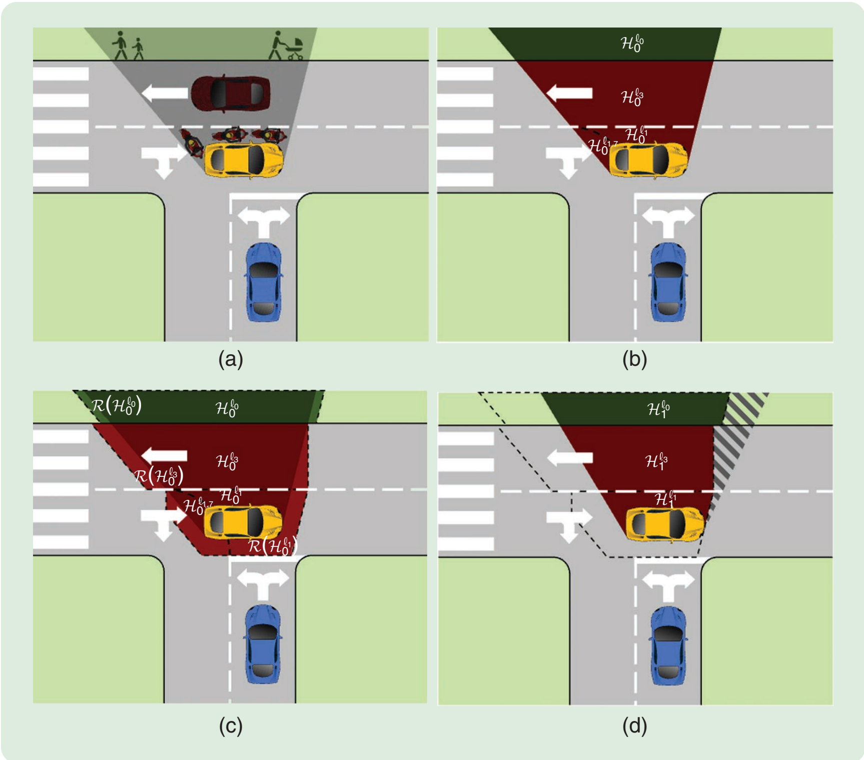

FIGURE 6. An automated (blue) ego vehicle is waiting at a T intersection. A (yellow) car passing by is occluding the field of view of the ego vehicle. Reachability analysis can be used to compute overapproximations of possible vehicle trajectories in the occluded region. (a) The occluded area could potentially hide vehicles or pedestrians. (b) The area can be represented as one or more sets. (c) The sets are propagated forward in time using reachability computations. (d) The sets reduce as information from new sensor readings is integrated. (Illustration from [6].)

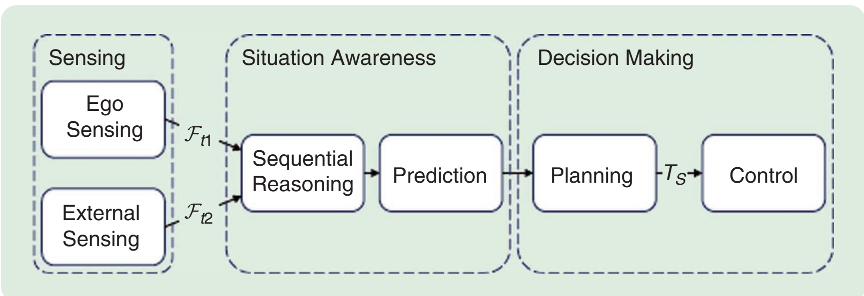

FIGURE 7. The planning and control of the ego vehicle are based on information received from the vehicle's own sensors together with external sensors, which form the situational awareness of the vehicle. (Illustration from [6].)

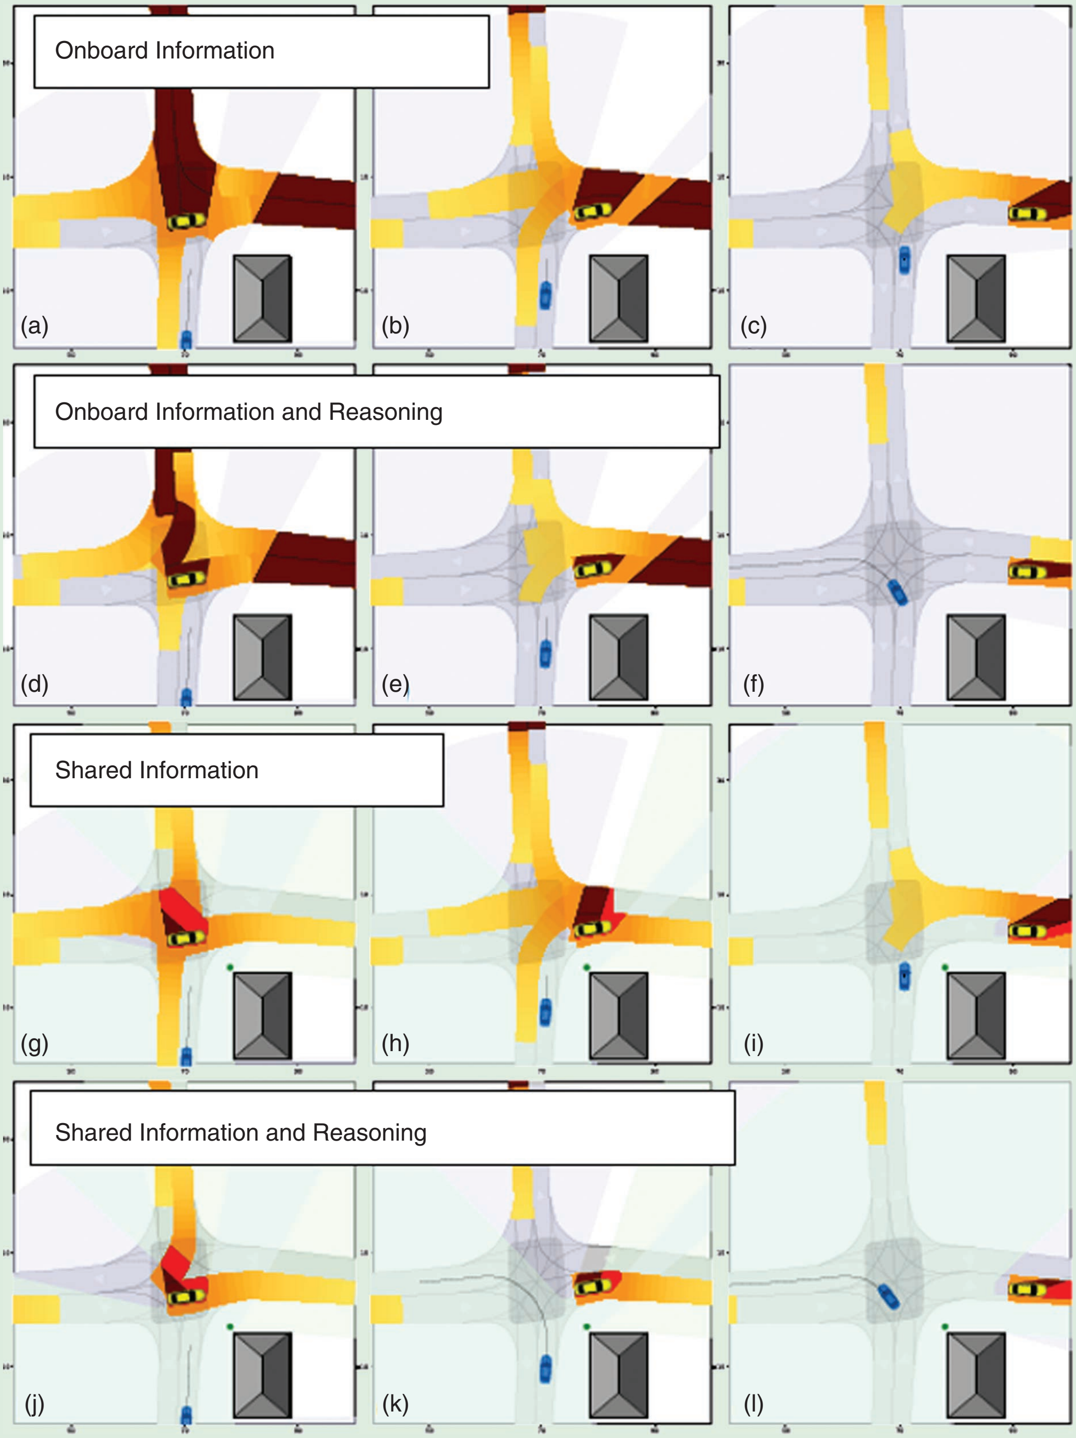

FIGURE 8. Simulated scenarios with the ego vehicle turning left at an intersection. Our research shows how using more sensor information and reasoning allows the vehicle to move faster through the intersection while still maintaining a guaranteed level of safety. (a)-(c) illustrate three snap shots when the ego vehicle uses only onboard information, while in (d)-(f) it uses onboard information together with reasoning based on reachability computations. In (g)-(i), the ego vehicle uses also shared information from a road side unit (blue dot by the building), and in (j)-(l) it uses the shared information together with reasoning. (Illustration from [6].)

Even 10–15 years ago, the leadership of some major corporations, such as Google and Tesla, considered autonomous driving to be a solved problem.

TRAFFIC CONTROL USING PHYSICS-INFORMED MACHINE LEARNING

Our second case study examines the application of machine learning techniques to address uncertainty in traffic control. This is important because poor traffic control results in an enormous waste of energy and time for road users worldwide. The sensors and actuators in this networked control system are connected and automated vehicles (CAVs). In particular, we consider the truck platooning technology we developed more than 10 years ago [10]. As trucks and truck platoons in Sweden maintain a speed of about 85 km/h, while cars drive at 110 or 120 km/h, cars overtake the trucks. In fact, a truck acts as a moving bottleneck, as there are fewer lanes for cars to use when they overtake the truck (Figure 10).

Lagrangian Traffic Control

Modern traffic control, based on the use of CAVs for mobile sensing and actuation, is sometimes referred to as "Lagrangian traffic control" due to its similarities to how Joseph-Louis Lagrange modeled fluid motion [11]. On the other hand, classical traffic control [12], [13], which is based on sensors fixed in road infrastructure to measure vehicle densities and flows, with control actuation implemented through digital speed signs, is denoted as "Eulerian traffic control" (see "Eulerian and Lagrangian Traffic Control"). In what follows, we outline a control architecture based on the Lagrangian approach (Figure 11) introduced recently (see, for example, [14], [15], [16], and [17]).

The objective is to reduce traffic congestion by smoothing the flow of vehicles. Flow and density measurements from a few connected probe vehicles (red) are used to learn a traffic model for the remaining vehicles (blue). One such classical model is the Lighthill–Whitham–Richards (LWR) partial differential equation (PDE)

$$\frac{\partial \rho(x,t)}{\partial t} + \frac{\partial V(\rho(x,t))\rho(x,t)}{\partial x} = 0$$

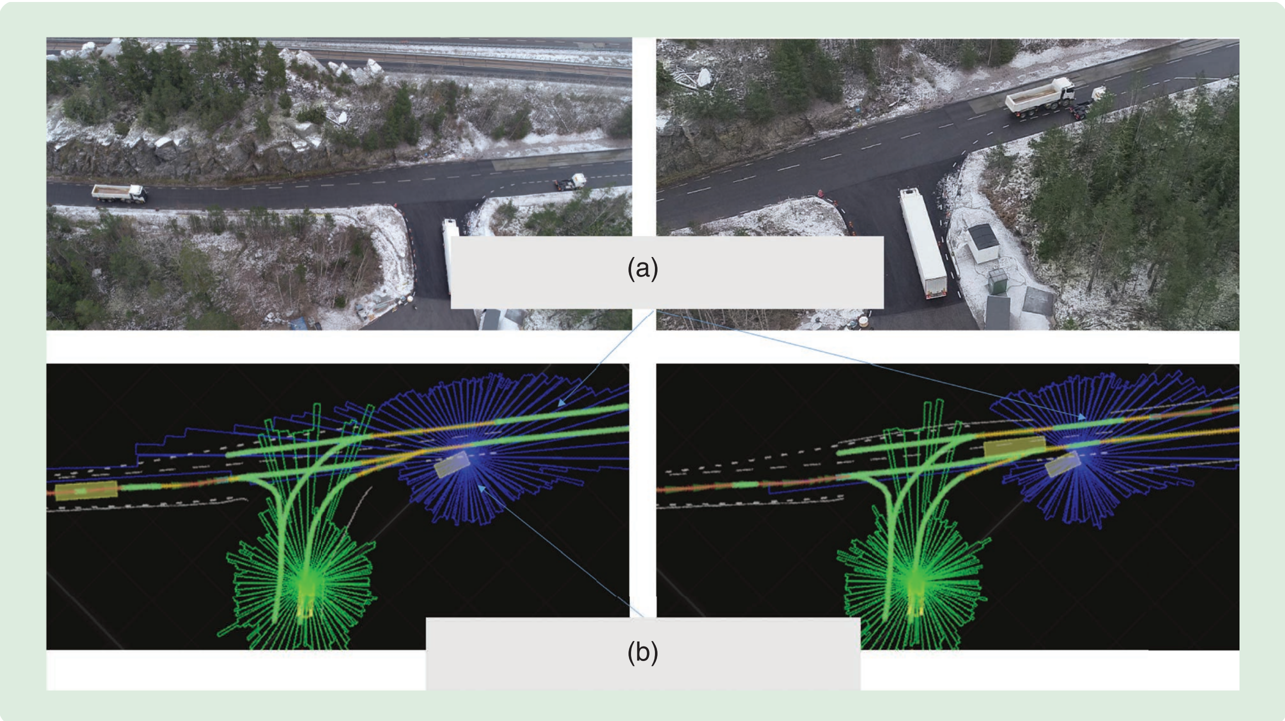

FIGURE 9. The experimental evaluation of automated safe driving under occlusion. An automated truck is waiting at a T intersection, processing sensor data to determine when it is safe to drive. (a) The ego vehicle can reason about whether the occluded area is occupied or not. (b) A roadside unit shares information about the occluded area with the ego vehicle. (Illustration from [6].)

Automated Truck Driving Under Low-Friction Conditions

Automated heavy-duty vehicles present a promising solution to improve the safety and efficiency of freight transport at scale. This emerging technology is increasingly relevant due to a lack of drivers, long-haul demands, and shorter delivery cycles. To ensure high operational reliability regardless of the diverse environment and changing road conditions, these vehicles must perform consistently not only for logistic continuity but also to ensure safety in scenarios that challenge professional human drivers.

Among the most critical scenarios for automated trucks are those involving snow and ice (see Figure S6), as they pose significant challenges to vehicle stability and control. Low-friction surfaces alter brake performance, steering authority, and lateral stability, particularly in articulated vehicle configurations, where trailer swing, traction loss, or jackknifing may occur. Under such conditions, robust and adaptive vehicle dynamics models become essential, especially when the vehicle's interaction with the road changes rapidly and cannot be precisely predicted.

Accurate and responsive modeling frameworks are thus fundamental for such real-time decision making under uncertainty. Recent advances in learning-based models capable of taking physical aspects into account have shown great potential [S12]. Integrating such models into the motion planning and control stack enables automated trucks to navigate low-friction terrains more safely and effectively, supporting critical maneuvers, such as precise stopping on ice, as described in Figure S7.

FIGURE S6 An automated truck operating under low-friction conditions.

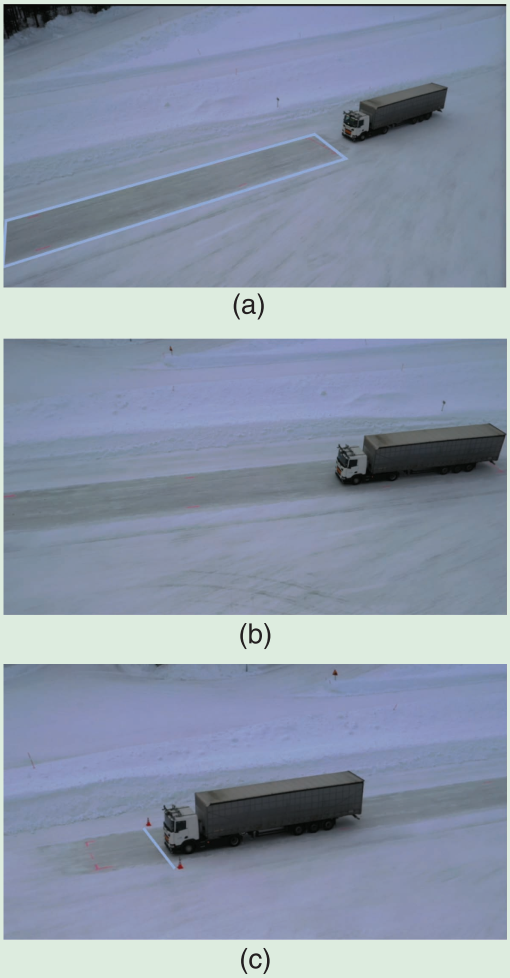

FIGURE S7 An automated truck executing a high-precision stopping maneuver on a low-friction surface (polished ice). The truck (a) approaches an ice patch, (b) slows down on the ice, and (c) stops at a designated point. The objective of the vehicle is to safely decelerate and stop at a cone-marked target location using only onboard sensors and a motion planning algorithm based on a physics-informed machine learning model.

REFERENCE

[S12] M. Selim, S. Bhat, and K. H. Johansson, "Motion planning using physics-informed LSTMs for autonomous driving," in *Proc. IEEE 27th Int. Conf. Intell. Transp. Syst. (ITSC)*, 2024, pp. 2251–2258, doi: 10.1109/ITSC58415.2024.10920067.

which describes how vehicle density $\rho = \rho(x, t)$ varies over space x and time t [18], [19]. Here, $\rho V(\rho)$ represents how traffic flow depends on density. The velocity of individual probe vehicles can be modeled by the ordinary differential equations (ODEs)

$$\dot{x}_i(t) = V(\rho(t, x_i(t))), \quad i = 1, \ldots, N$$

where N is the number of probe vehicles considered in a given scenario. The traffic system to be controlled is consequently governed by a coupled system of PDEs and

The crash rate per million miles for automated vehicles remains more than 10 times higher than that of human-driven vehicles.

ODEs, which pose significant mathematical challenges [20], [21].

The traffic model learned from vehicle data is used to reconstruct the traffic state throughout the considered road segment, as well as for parts where the probe vehicles have not been driving. The model and the reconstructed state form the basis for predicting the traffic state. Similar to the above traffic scenario, this can involve predicting the evolution of a high-density area downstream. Based on the prediction, the controller computes the receding horizon control actions for the trucks or truck platoons to mitigate congestion by potentially slowing down some of the trucks slightly.

Learning-Based Traffic State Estimation

The traffic model is learned using physics-informed neural networks [23], as proposed in [16] and [24]. We modify the original LWR model by adding the second-order term $\gamma^2 \partial^2 \rho / \partial x^2$ to the right-hand side, obtaining a "diffusively corrected" LWR model. This model is learned using sampled measurements $\{x_i(t), \rho_i(t), V_i(t)\}$ from probe vehicles $i = 1, \ldots, N$. The goal is to obtain a model for the optimal density $\hat{\rho}^\star = \operatorname{argmin}_{\hat{\rho}} \int_0^T \| \rho(t, \cdot) - \hat{\rho}(t, \cdot) \|^2 dt$. We do that by using feedforward neural networks to represent the maps $\hat{\rho}_{\Theta_\rho} = \hat{\rho}_{\Theta_\rho}(t,x)$ and $\hat{V}_{\Theta_V} = \hat{V}_{\Theta_V}(\rho)$, where Θ_ρ and Θ_V denote the parameters of the networks. To train the neural networks, the following optimization problem is considered:

$$
\begin{aligned}
\min_{\Theta_\rho, \Theta_V} \; & \frac{1}{N} \sum_{i=1}^{N} \int_0^T \Big\{ \big| \hat{\rho}_{\Theta_\rho}(t, x_i(t)) - \rho_i(t) \big|^2 \\
& \qquad + \big| \hat{V}_{\Theta_V}(\rho_i(t)) - V_i(t) \big|^2 \Big\} dt \\
\text{subject to} \; & \frac{\partial \hat{\rho}_{\Theta_\rho}}{\partial t} + \frac{\partial \hat{V}_{\Theta_V}(\hat{\rho}_{\Theta_\rho}) \hat{\rho}_{\Theta_\rho}}{\partial x} = \gamma^2 \frac{\partial^2 \hat{\rho}_{\Theta_\rho}}{\partial x^2} \\
& \left| \frac{\partial \hat{V}_{\Theta_V}}{\partial \rho}(\rho) \right|_+ = 0
\end{aligned}
$$

where the constraints impose the desired physical properties of the traffic model. Through Lagrangian relaxation, we obtain the optimization problem

$$
\begin{aligned}
\min_{\Theta_\rho, \Theta_V} \max_{\lambda_\rho, \lambda_V} \; & \frac{1}{N} \sum_{i=1}^{N} \int_0^T \Big\{ \big| \hat{\rho}_{\Theta_\rho}(t, x_i(t)) - \rho_i(t) \big|^2 \\
& \qquad + \big| \hat{V}_{\Theta_V}(\rho_i(t)) - V_i(t) \big|^2 \Big\} dt \\
& + \lambda_\rho \iint_{[0,T]\times[0,L]} \bigg| \frac{\partial \hat{\rho}_{\Theta_\rho}}{\partial t}(\nu) + \frac{\partial \hat{V}_{\Theta_V}(\hat{\rho}_{\Theta_\rho}) \hat{\rho}_{\Theta_\rho}}{\partial x}(\nu) \\
& \qquad - \gamma^2 \frac{\partial^2 \hat{\rho}_{\Theta_\rho}}{\partial x^2}(\nu) \bigg| d\nu \\
& + \lambda_V \int_0^1 \left| \frac{\partial \hat{V}_{\Theta_V}}{\partial \rho}(\rho) \right|_+^2 d\rho \\
=: & \min_{\Theta_\rho, \Theta_V} \max_{\lambda_\rho, \lambda_V} \mathcal{L}_{\lambda_\rho, \lambda_V}(\Theta_\rho, \Theta_V).
\end{aligned}
$$

FIGURE 10. When (a) a three-vehicle truck platoon drives on (b) a highway, it acts as (c) a moving bottleneck with respect to the rest of the traffic. (Source: Scania; used with permission.)

Eulerian and Lagrangian Traffic Control

Traffic control has traditionally been compared to the control of a fluid [S13]. The Swiss mathematician and physicist Leonhard Euler (1707–1783) focused on specific locations in the space through which fluid flows as time passes when he developed his mathematical model of fluid motion [S14]. Inspired by this view, common traffic control based on sensors and actuators fixed in infrastructure is referred to as "Eulerian traffic control" [S15], [S16], [S17]. As illustrated in Figure S8, such sensors can be radar or camera sensors located above a highway that measure the density and flow of vehicles at that specific location. Actuators can be digital road signs that indicate the desired speed that drivers should follow.

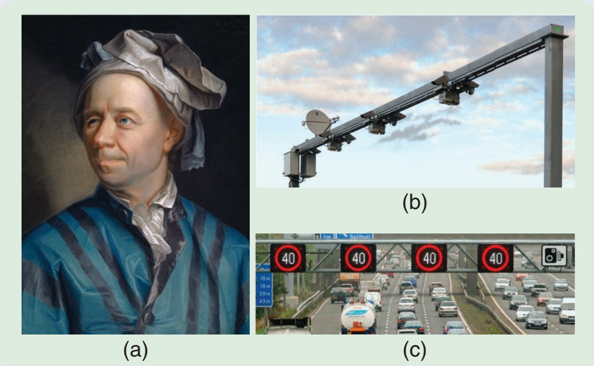

FIGURE S8 (a) Leonhard Euler has his name attached to traffic control with (b) fixed sensors and (c) actuators.

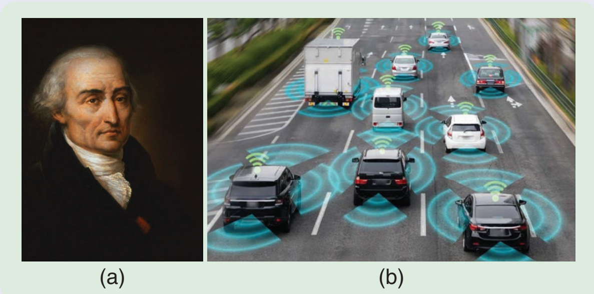

FIGURE S9 (a) Joseph-Louis Lagrange. (b) Lagrangian traffic control is based on mobile sensors and actuators, such as traffic state information from wirelessly connected probe vehicles and controlled automated vehicles, respectively.

The Italian–French mathematician and physicist Joseph-Louis Lagrange (1736–1813) considered fluid motion from the viewpoint that the observer follows an individual fluid parcel as it moves through space and time [11]. The corresponding traffic control paradigm is based on mobile sensors and actuators, as in Figure S9. Probe vehicles act as mobile sensors by estimating and communicating the traffic state information they observe [S18]. Some CAVs are controlled (for example, it could be the truck in the figure), and thus, they serve as actuators in vehicle flow [S19].

Eulerian traffic control relies on a substantial fixed infrastructure and therefore has a high deployment cost and limited flexibility. Emerging Lagrangian traffic control has great potential; however, a system-theoretic foundation is needed for such control systems [S20]. Some data-based methods for learning Lagrangian traffic models and predicting their traffic state are discussed in this article and in further detail in [S21].

REFERENCES

[S13] L. C. Edie, "Discussion of traffic stream measurements and definitions," in *Proc. 2nd Int. Symp. Theory Traffic Flow*, Paris, France: OECD, 1963, pp. 139–154.

[S14] L. Euler, "Principes généraux du mouvement des fluides," *Mémoires de L'Académie Des Sciences de Berlin*, vol. 11, pp. 274–315, 1757.

[S15] A. Ferrara, S. Sacone, and S. Siri, *Freeway Traffic Modeling and Control*. Cham, Switzerland: Springer-Verlag, 2018.

[S16] M. Papageorgiou, H. Hadj-Salem, and J.-M. Blosseville, "Modelling and real-time control of traffic flow on the southern part of Boulevard Peripherique in Paris: Part I: Modelling," *Transp. Res. Part A, General*, vol. 24, no. 5, pp. 345–359, 1991, doi: 10.1016/0191-2607(90)90047-A.

[S17] A. Hegyi, B. De Schutter, and H. Hellendoorn, "Model predictive control for optimal coordination of ramp metering and variable speed limits," *Transp. Res. Part C: Emerg. Technol.*, vol. 13, no. 3, pp. 185–209, Jun. 2005, doi: 10.1016/j.trc.2004.08.001.

[S18] D. B. Work, O. -P. Tossavainen, Q. Jacobson, and A. M. Bayen, "Lagrangian sensing: Traffic estimation with mobile devices," in *Proc. Amer. Control Conf.*, St. Louis, MO, USA, 2008, pp. 1536–1543, doi: 10.1109/ACC.2009.5160332.

[S19] M. Čičić and K. H. Johansson, "Traffic regulation via individually controlled automated vehicles: A cell transmission model approach," in *Proc. 21st IEEE Int. Conf. Intell. Transp. Syst. (ITSC)*, 2018, pp. 766–771, doi: 10.1109/ITSC.2018.8569960.

[S20] M. L. Delle Monache, T. Liard, B. Piccoli, R. Stern, and D. Work, "Traffic reconstruction using autonomous vehicles," *SIAM J. Appl. Math.*, vol. 79, no. 5, pp. 1748–1767, 2019, doi: 10.1137/18M1217000.

[S21] M. Barreau, M. Aguiar, J. Liu, and K. H. Johansson, "Physics-informed learning for identification and state reconstruction of traffic density," in *Proc. 60th IEEE Conf. Decis. Control*, 2021, pp. 2653–2658, doi: 10.1109/CDC45484.2021.9683295.

This problem can be solved numerically by primal–dual gradient descent iterations based on the primal problem

$$\Theta_\rho^*(\lambda_\rho),\, \Theta_V^*(\lambda_V) = \arg\min_{\Theta_\rho,\Theta_V} \mathcal{L}_{\lambda_\rho,\lambda_V}(\Theta_\rho,\, \Theta_V)$$

and the dual problem

$$\lambda_\rho^*(\Theta_\rho), \lambda_V^*(\Theta_V) = \arg\max_{\lambda_\rho,\lambda_V} \mathcal{L}_{\lambda_\rho,\lambda_V}(\Theta_\rho,\Theta_V).$$

See [16] for further details on the implementation.

We simulate in SUMO [25] a 2.5-km road segment over 7 min, where the density ρ is illustrated by color (Figure 12). The density of vehicles varies between high (red) and low (blue), with black curves indicating how individual probe vehicles drive the road segment. Note that when the density is lower, the probe vehicles drive faster, and when they experience congestion, they drive slower, as expected

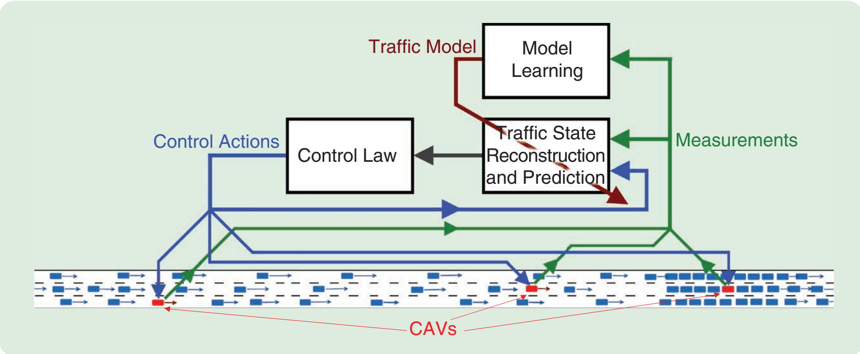

FIGURE 11. The Lagrangian traffic control architecture. CAVs are used to gather traffic state information and to implement control actions. The controller is based on a machine learning model that enables the prediction of future traffic states, which are utilized by the control law. (Illustration adopted from [22].)

domain, the estimated traffic density is quite close to the true one.

Vehicle Control to Dissipate Traffic Congestion

Vehicles moving slower than the average traffic act as moving bottlenecks, slowing other vehicles [14], [15], [21], [26]. By actively controlling trucks or truck platoons, it is possible to slightly reduce or increase the traffic flow and therefore influence downstream traffic congestion (Figure 13). The figure illustrates a situation where a car accident happens at 50 km at time t_0. The accident gives rise to high density (the yellow area), propagating backward in space. The red curve indicates that a truck platoon is driving at a constant speed (the straight line) until it reaches the congested area (yellow)

from the traffic model. Based on samples obtained by the probe vehicles, the reconstructed traffic state is shown using the physics-informed machine learning approach. Despite the sparse sampling of the spatial–temporal

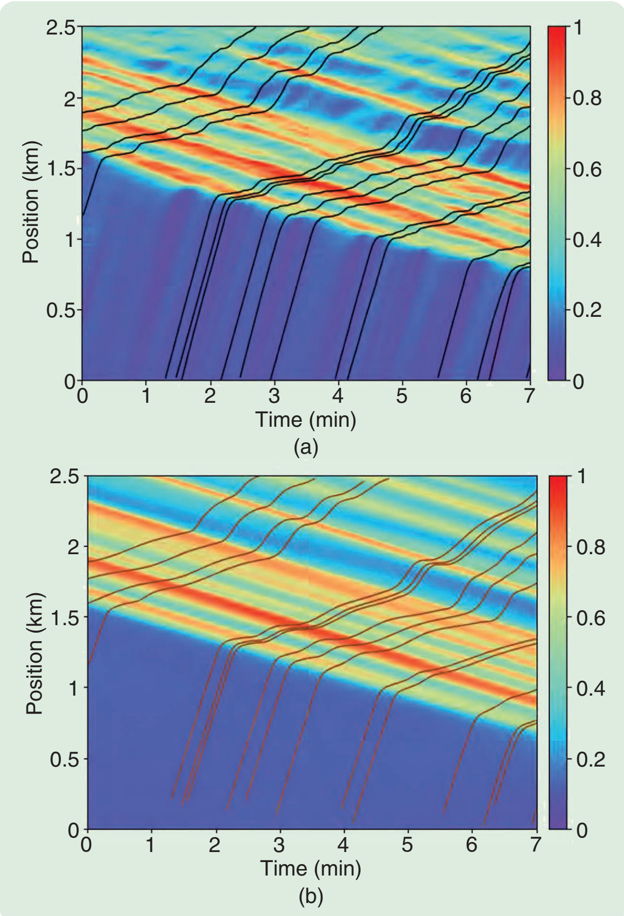

FIGURE 12. (a) Simulated traffic over a 2.5-km road segment. The color indicates vehicle density; for instance, the red areas have high density. The black curves are the trajectories of probe vehicles. (b) Estimated densities based on collected data using the described physics-informed machine learning approach. The estimated density agrees quite well with the true density. (Numerical results of [16].)

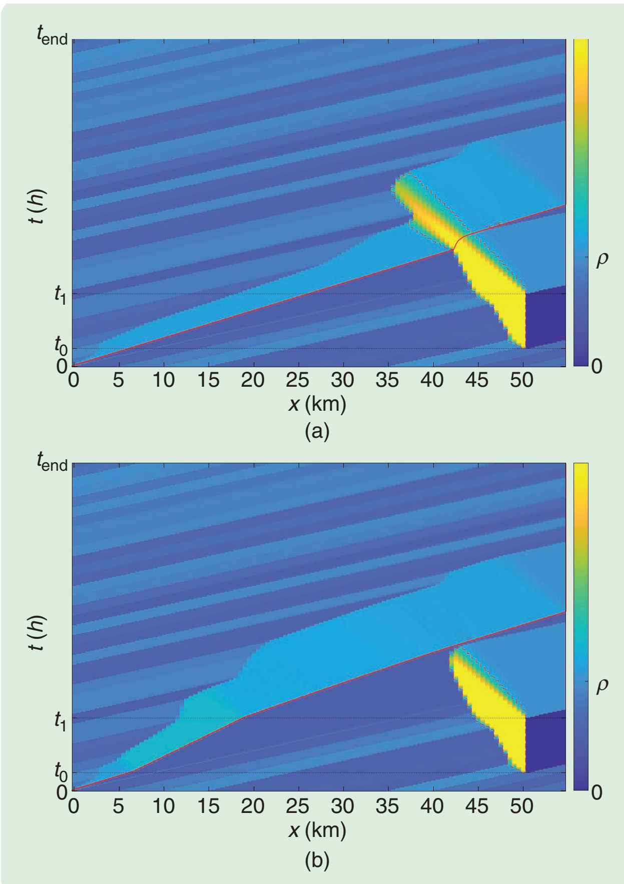

FIGURE 13. (a) A truck platoon (the red curve) driving through a highly congested area indicated by the yellow region. (b) By controlling the velocity of the platoon, it is possible to reduce the traffic flow into the congested area and thereby mitigate the influence of the congestion on the total travel time for all the vehicles driving through the considered road segment. (Numerical results of [14].)

when it slows down. After the congested area, it continues with a higher constant speed again. The size of the congested area corresponds to the severity of the accident in the sense of how much it influences the travel time for general traffic. Using the estimated and predicted traffic status, it is possible to slow down the platoon to reduce the congested area. This slowing down is illustrated in the figure. Note that the slope of the red curve is slightly larger with platoon control compared to without. The slowdown of the platoon leads to the density behind it being slightly higher, and as a consequence, the (yellow) congested area is reduced. It can be shown that the overall travel time for all cars in this scenario is significantly reduced due to the platoon control [14], [27].

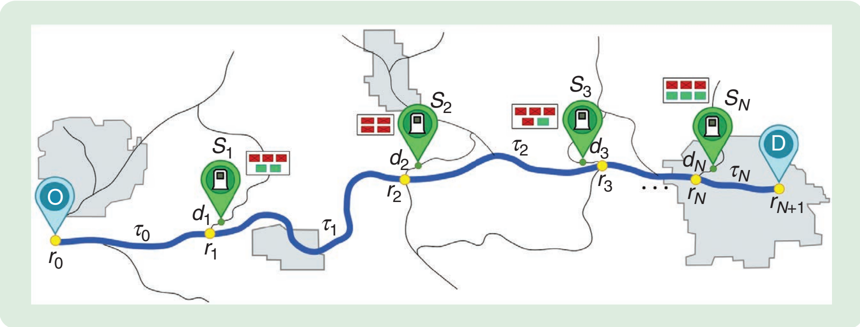

FIGURE 14. The route model of each truck, where the preplanned route between the origin and destination is represented by the blue path. Charging stations are shown by green labels. Green and red blocks at each station indicate the availability of the charging ports. (Illustration from [29].)

ROLLOUT-BASED PLANNING FOR ELECTRIC TRUCK CHARGING COORDINATION

The third case study that we consider to mitigate uncertainty in transport systems is on electric truck charging coordination for large heavy-duty vehicle fleets. In many regions of the world, there is an electric charging infrastructure for cars. We do not yet have such an infrastructure to charge heavy vehicles, as they require significantly more power to charge, and the number of electric heavy vehicles is still relatively small. The development of such a charging infrastructure poses some interesting control problems under uncertainty, which we describe next [28].

Consider an example with 1,000 electric trucks that travel daily on the Swedish road network. We are interested in the following problem: How can electric trucks decide where and when to charge, given that they have preplanned routes, there is limited charging capacity, and travel time and energy consumption are uncertain (Figure 14)? The solution to the problem should be scalable in the sense that each vehicle computes its own charging plan, and the solution should acknowledge privacy constraints so that trucks never have to reveal their travel plans to other trucks.

A distributed rollout-based solution [30] is developed for truck planning. It is based on waiting time forecast models for charging stations. Each truck interacts with charging stations along its route to obtain estimated waiting times, which are used to simulate the worst-case scenario if a truck charges at the next station or if it does not, following the core idea of rollout-based planning. An optimal plan for a single truck is illustrated together with its corresponding waiting times (Figure 15). The underlying data are obtained from realistic freight transport missions [28]. See "Rollout in Approximate DP and Reinforcement Learning" for an introduction to rollout-based planning.

Calculating the average total waiting time for all 1,000 electric trucks per day (Figure 16) shows that the

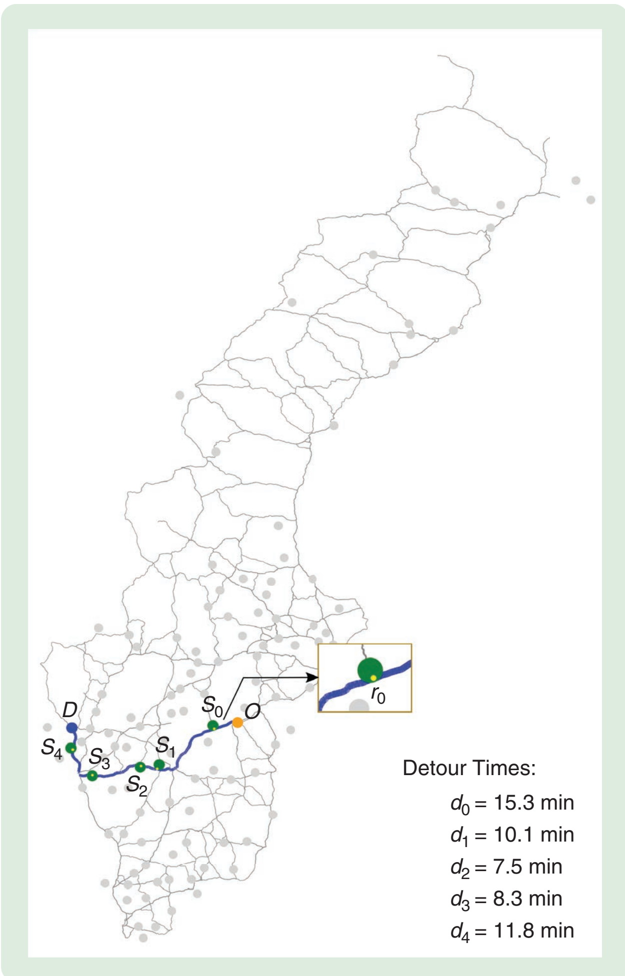

FIGURE 15. A charging plan for a single truck going from its origin to its destination in the southern part of Sweden. The truck charges five times. (Illustration from [28].)

Rollout in Approximate DP and Reinforcement Learning

Rollout is an approximate solution method for solving optimal control problems with a finite or infinite horizon. The main idea of rollout algorithms, which involves obtaining an improved policy starting from some suboptimal policy, has been explored in several DP and reinforcement learning settings, including policy iteration and computer games, for example, in [S22] and [S23]. The term "rollout" was coined in [S23] in the context of backgammon, where the name refers to rolling the dice. Since then, rollout and its variants have been applied to combinatorial optimization [S24], [S25], [S26], stochastic scheduling [S27], vehicle routing [S28], Bayesian optimization [S29], [S30], multiagent problems [S31], and the AlphaGo program [S32]; see the monograph [S33] for a comprehensive introduction to rollout and more references. Although rollout was initially designed for problems with discrete control space, classical model predictive control (MPC) with continuous control space can also be viewed as a form of rollout [S34], [S35], [S36]. In what follows, we focus on the finite-horizon stochastic optimal control problem and introduce the rollout, along with some of its variants. Most of the material is adapted from [S33] and [S37].

FINITE-HORIZON STOCHASTIC OPTIMAL CONTROL

A finite-horizon stochastic optimal control problem involves a state equation

$$x_{k+1} = f_k(x_k, u_k, w_k), \ k = 0, 1, \ldots, N-1 \tag{S1}$$

where x_k, u_k, and w_k are the state, control, and disturbance, respectively, at stage k. The state x_k takes values in the state space X_k, and the control is taken from the control constraint set $U_k(x_k)$ that may depend on the state. The probability distribution of w_k depends on (x_k, u_k). A policy π is a sequence of functions $\{\mu_0, \mu_1, \ldots, \mu_{N-1}\}$, with $\mu_k(x_k) \in U_k(x_k)$. We associate each policy π with a cost function J_π, defined pointwise for each x_0 as

$$J_\pi(x_0) = \underset{\substack{w_k \\ k=0,\ldots,N-1}}{E}\left\{g_N(x_N) + \sum_{k=0}^{N-1} g_k(x_k, \mu_k(x_k), w_k)\right\}$$

subject to constraints imposed by the state equation (S1), where $E\{\cdot\}$ denotes expectation, $g_k(x_k, u_k, w_k)$ is the cost at stage k, and $g_N(x_N)$ is the terminal cost at stage N. Our goal is to compute the optimal cost function

$$J^*(x_0) = \min_\pi J_\pi(x_0), \text{ for all } x_0.$$

Although the optimal cost J^* can be computed via DP in principle, it is not tractable for large state spaces. Rollout comes in handy for providing an approximate solution with a reasonable computational load.

EXACT ROLLOUT AND PERFORMANCE GUARANTEE

Given a policy $\pi = \{\mu_0, \ldots, \mu_{N-1}\}$ computed offline, which is called the *base policy*, rollout computes a policy through online optimization. The policy obtained is known as the *rollout policy* and is denoted as $\tilde{\pi} = \{\tilde{\mu}_0, \ldots, \tilde{\mu}_{N-1}\}$. In particular, upon reaching state x_k, the control $\tilde{\mu}_k(x_k)$ is computed via

$$\tilde{\mu}_k(x_k) \in \arg\min_{u_k \in U_k(x_k)} \underset{w_k}{E}\{g_k(x_k, u_k, w_k) + J_{k+1,\pi}(x_{k+1})\} \tag{S2}$$

subject to (S1), and for $k = 1, \ldots, N-1$,

$$J_{k,\pi}(x_k) = \underset{\substack{w_\ell \\ \ell=k,\ldots,N-1}}{E}\left\{g_N(x_N) + \sum_{\ell=k}^{N-1} g_\ell(x_\ell, \mu_\ell(x_\ell), w_\ell)\right\}$$

and $J_{N,\pi}(x_N) = g_N(x_N)$. The values $J_{k+1,\pi}(x_{k+1})$ in (S2) can only in rare cases be derived in closed form. Instead, they are computed online via simulation. An illustration of such a rollout to compute $J_{k+1,\pi}(x_{k+1})$ via simulation applied to backgammon is presented in Figure S10.

In the literature, the minimization (S2), which involves only the control of the present stage, is called *one-step* look ahead. In *ℓ-step look ahead*, the control of rollout policy $\tilde{\mu}_k(x_k)$ is computed by minimizing over the control of the current stage as well as over the functions $\mu'_{k+1}, \ldots, \mu'_{k+\ell-1}$, with future costs represented by the function $J_{k+\ell,\pi}$. In other words, the online computation of rollout with ℓ-step look ahead solves an ℓ-stage optimal control problem. However, once the optimum is achieved at $\hat{u}_k, \hat{\mu}_{k+1}, \ldots, \hat{\mu}_{k+\ell}$, only the first control $\hat{u}_k$ is used to define the rollout policy at the stage k, $\tilde{\mu}_k(x_k) = \hat{u}_k$, while the functions $\hat{\mu}_{k+1}, \ldots, \hat{\mu}_{k+\ell}$ are discarded.

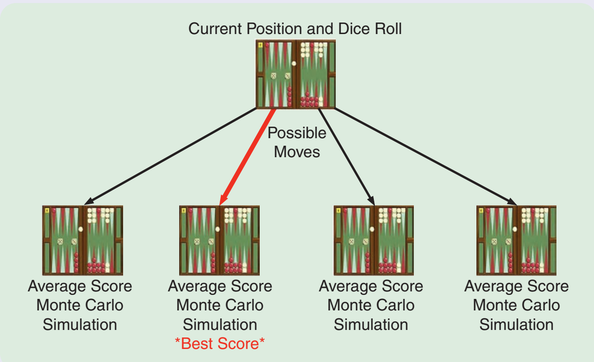

FIGURE S10 The rollout for backgammon. At a given position and roll of the dice, the set of all possible moves is generated, and the outcome of the game for each move is evaluated by "rolling out" (simulating to the end) many games using a suboptimal or heuristic backgammon player and by Monte Carlo averaging the scores. The move that results in the best average score is selected for play. (Source: Athena Scientific [S37]; used with permission.)

Regardless of the number of look-ahead steps, it can be shown that the rollout policy is guaranteed to outperform the base policy in the sense that

$$J_{k,\tilde{\pi}}(x_k) \le J_{k,\pi}(x_k), \text{ for all } x_k \in X_k,\ k = 0, 1, \ldots, N. \tag{S3}$$

In particular, $J_{\tilde{\pi}}(x_0) = J_{0,\tilde{\pi}}(x_0) \le J_{0,\pi}(x_0) = J_\pi(x_0)$.

VARIANTS OF ROLLOUT

One challenge of rollout is to compute the values $J_{k,\pi}(x_k)$ online. A popular variant that aims to address this difficulty is *truncated rollout*, where the simulation used to estimate the values runs for m stages. Some offline computed functions $\tilde{J}_{k+m+1}$ are then used to capture the cost beyond these stages. In particular, for rollout with one-step look-ahead minimization (S2), the values $J_{k+1,\pi}(x_{k+1})$ are replaced by

$$\underset{\substack{w_\ell \\ \ell=k+1,\ldots,k+m}}{E}\left\{\tilde{J}_{k+m+1}(x_{k+m+1}) + \sum_{\ell=k+1}^{k+m} g_\ell(x_\ell,\ \mu_\ell(x_\ell),\ w_\ell)\right\}.$$

See Figure S11.

Alternatively, we may fix the disturbances $w_{k+1}, \ldots, w_{N-1}$ at some typical values $\bar{w}_{k+1}, \ldots, \bar{w}_{N-1}$ and replace $J_{k+1,\pi}(x_{k+1})$ by

$$g_N(x_N) + \sum_{\ell=k}^{N-1} g_\ell(x_\ell,\ \mu_\ell(x_\ell),\ \bar{w}_\ell).$$

In this case, only one sample trajectory is needed for each x_{k+1}. This is called *certainty equivalence*, and it was introduced for rollout in [S27]. For both of these rollout variants, the theoretical guarantee for exact rollout in (S3) no longer holds. However, extensive computational studies have shown that rollout policies obtained via these variants often outperform base policies by large margins. This is due to the connection to Newton's method, which holds even when approximations are involved; see [30].

Another computational challenge of rollout is the minimization in (S2). For problems where the minimum cannot be attained exactly, one may replace the control constraint set $U_k(x_k)$ with some suitably constructed subset $\bar{U}_k(x_k) \subset U_k(x_k)$ such that $\mu_k(x_k) \in \bar{U}_k(x_k)$ and compute the corresponding minimum instead. This is called *simplified rollout*. It is related to the notion in MPC that obtaining a feasible solution is sufficient [S38]. A special form of this variant is *multiagent rollout* tailored for multiagent systems. For this variant, the performance guarantee given in (S3) remains intact.

If there are sufficient computational resources, the rollout policy defined by the minimization (S2) can also be enhanced. One such variant is *parallel rollout* (also known as *rollout with multiple heuristics*), introduced in [S24]. This scheme assumes that there are multiple base policies $\pi^1, \pi^2, \ldots, \pi^m$ and that we replace the function $J_{k+1,\pi}$ with $\hat{J}_{k+1}$, which is defined pointwise as

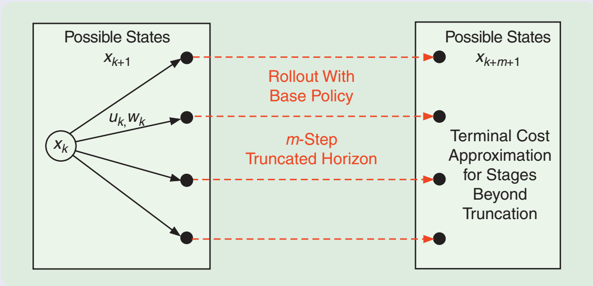

FIGURE S11 Truncated rollout. One-step look ahead is followed by a simulation of the base policy for m steps. An approximate remaining cost $\tilde{J}_{k+m+1}(x_{k+m+1})$ is then added to the cost of the simulation, which depends on the state x_{k+m+1} obtained at the end of the rollout. (Source: Athena Scientific [S37]; used with permission.)

$$\hat{J}_{k+1}(x_{k+1}) = \min\{J_{k+1,\pi^1}(x_{k+1}), \ldots, J_{k+1,\pi^m}(x_{k+1})\}.$$

It can be shown that for $\ell = 1, 2, \ldots, m$,

$$J_{k,\tilde{\pi}}(x_k) \le J_{k,\pi^\ell}(x_k), \text{ for all } x_k \in X_k,\ k = 0, 1, \ldots, N.$$

This variant has been applied for partial state information problems [S39] as well as problems with continuous state and control spaces [S40], [S41], [S42].

REFERENCES

[S22] B. Abramson, "Expected-outcome: A general model of static evaluation," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 12, no. 2, pp. 182–193, Feb. 1990, doi: 10.1109/34.44404.

[S23] G. Tesauro and G. R. Galperin, "On-line policy improvement using Monte-Carlo search," in *Proc. 10th Int. Conf. Neural Inf. Process. Syst. (NIPS)*, 1996, pp. 1068–1074.

[S24] D. P. Bertsekas, J. N. Tsitsiklis, and C. Wu, "Rollout algorithms for combinatorial optimization," *J. Heuristics*, vol. 3, no. 3, pp. 245–262, 1997, doi: 10.1023/A:1009635226865.

[S25] D. Bertsimas and R. Demir, "An approximate dynamic programming approach to multidimensional knapsack problems," *Manage. Sci.*, vol. 48, no. 4, pp. 550–565, 2002, doi: 10.1287/mnsc.48.4.550.208.

[S26] Y. Li and D. P. Bertsekas, "Most likely sequence generation for *n*-grams, transformers, HMMs, and Markov chains, by using rollout algorithms," 2024, *arXiv:2403.15465*.

[S27] D. P. Bertsekas and D. A. Castanon, "Rollout algorithms for stochastic scheduling problems," *J. Heuristics*, vol. 5, no. 1, pp. 89–108, 1999, doi: 10.1023/A:1009634810396.

[S28] N. Secomandi, "A rollout policy for the vehicle routing problem with stochastic demands," *Oper. Res.*, vol. 49, no. 5, pp. 796–802, 2001, doi: 10.1287/opre.49.5.796.10608.

[S29] R. R. Lam, K. E. Willcox, and D. H. Wolpert, "Bayesian optimization with a finite budget: An approximate dynamic programming approach," in *Proc. 30th Int. Conf. Neural Inf. Process. Syst.*, 2016, pp. 883–891.

[S30] D. Bertsekas, "Rollout algorithms and approximate dynamic programming for bayesian optimization and sequential estimation," 2022, *arXiv:2212.07998*.

[S31] D. P. Bertsekas, "Multiagent reinforcement learning: Rollout and policy iteration," *IEEE/CAA J. Autom. Sin.*, vol. 8, no. 2, pp. 249–272, Feb. 2021, doi: 10.1109/JAS.2021.1003814.

[S32] D. Silver et al., "Mastering the game of Go with deep neural networks and tree search," *Nature*, vol. 529, no. 7587, pp. 484–489, 2016, doi: 10.1038/nature16961.

(*Continued*)

Rollout in Approximate DP and Reinforcement Learning (*Continued*)

[S33] D. P. Bertsekas, *Rollout, Policy Iteration, and Distributed Reinforcement Learning*. Belmont, MA, USA: Athena Scientific, 2020.

[S34] D. P. Bertsekas, "Dynamic programming and suboptimal control: A survey from ADP to MPC," *Eur. J. Control*, vol. 11, nos. 4–5, pp. 310–334, 2005, doi: 10.3166/ejc.11.310-334.

[S35] Y. Li, "Approximate methods of optimal control via dynamic programming models," Ph.D. thesis, Kungliga Tekniska högskolan, Stockholm, Sweden, 2023.

[S36] D. P. Bertsekas, "Model predictive control and reinforcement learning: A unified framework based on dynamic programming," *IFAC-PapersOnLine*, vol. 58, no. 18, pp. 363–383, 2024, doi: 10.1016/j.ifacol.2024.09.056.

[S37] D. P. Bertsekas, *A Course in Reinforcement Learning*, 2nd ed. Belmont, MA, USA: Athena Scientific, 2024.

[S38] P. O. M. Scokaert, D. Q. Mayne, and J. B. Rawlings, "Suboptimal model predictive control (feasibility implies stability)," *IEEE Trans. Autom. Control*, vol. 44, no. 3, pp. 648–654, Mar. 1999, doi: 10.1109/9.751369.

[S39] H. S. Chang, R. Givan, and E. K. P. Chong, "Parallel rollout for online solution of partially observable Markov decision processes," *Discrete Event Dyn. Syst.*, vol. 14, no. 3, pp. 309–341, 2004, doi: 10.1023/B:DISC.0000028199.78776.c4.

[S40] U. Rosolia and F. Borrelli, "Learning model predictive control for iterative tasks. a data-driven control framework," *IEEE Trans. Autom. Control*, vol. 63, no. 7, pp. 1883–1896, Jul. 2018, doi: 10.1109/TAC.2017.2753460.

[S41] Y. Li, K. H. Johansson, J. Mårtensson, and D. P. Bertsekas, "Data-driven rollout for deterministic optimal control," in *Proc. 60th IEEE Conf. Decis. Control (CDC)*, 2021, pp. 2169–2176, doi: 10.1109/CDC45484.2021.9683499.

[S42] Y. Li, A. Karapetyan, N. Schmid, J. Lygeros, K. H. Johansson, and J. Mårtensson, "Parallel model predictive control for deterministic systems," 2023, *arXiv:2309.14560*.

proposed rollout-based solution gives approximately 50% shorter waiting times compared to an offline strategy. See "Multifleet Platoon Coordination for Large-Scale Freight Transportation Systems" for another truck planning problem.

CONCLUSIONS

Let us conclude this lecture with some perspectives. The future is not just more data and more machine learning computations. The future of advancing intelligent transportation consists of a systematic and effective mitigation of uncertainty. The control community has, for decades, been developing tools to mitigate uncertainty. In the era of big data and artificial intelligence, there is a real need not only for these tools but also for the development of new and advanced architectures, models, and algorithms for accommodating uncertainty. In this lecture, we illustrated, through three case studies, how uncertainty can be represented and compensated for using safety-first reachable set computations, a mix of physics- and data-driven models, and stochastic forecasting with distributed rollout-based solutions, which can be viewed as a first step toward this future.

Finally, I would like to congratulate all students in the audience; you have chosen the right scientific discipline. This scientific discipline of *control systems* is where significant developments are currently taking place and will continue to do so in the future. The control community not only possesses the expertise in creating and deploying the mathematical and computational tools but also the right spirit to address the great challenges of our generation.

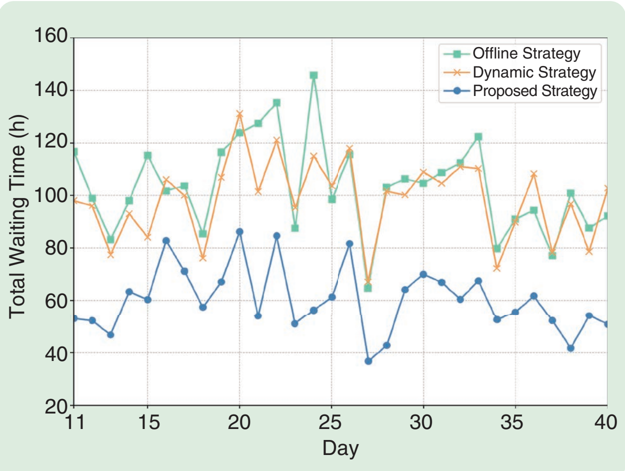

FIGURE 16. The average total waiting time for all 1,000 electric trucks per day over about a month. The rollout-based solution is significantly better at utilizing existing and uncertain resources. (Illustration from [29].)

THANKS

I would like to thank the audience for your attention. I also extend a special thanks to many collaborators (Ph.D. students, postdoctoral fellows, research colleagues, and industrial engineers) who have developed the research results that I have presented. Special thanks for all the help and feedback I received in preparing for this presentation go to Matthieu Barreau, Mladen Čičić, Frank Jiang, Ting Bai, Yuchao Li, Truls Nyberg, Vandana Narri, Mahmoud Selim, Miguel Aguiar, Jonas Mårtensson, Henrik Sandberg, Dimos Dimarogonas, Dan Work, and Carlos Canudas-de-Wit. I would like to extend a special thank you to Ting Bai, Dimitri Bertsekas, Mladen Čičić, Frank Jiang, Yuchao Li, and Mahmoud Selim for their help in preparing the sidebars of this article. I am grateful for generous and long-term financial support from the Knut and Alice Wallenberg Foundation, the Swedish Research Council, the Swedish Strategic Research Foundation, the Swedish Agency for Innovation Systems, the European Union's Funding Program

Multifleet Platoon Coordination for Large-Scale Freight Transportation Systems

Hub-based platooning enables trucks to form platoons at designated points of interest, supporting on-time deliveries, reducing fuel consumption, lowering operational costs, and improving driving safety [10]; see Figure S12. Coordinating platoons across multiple fleets is crucial to fully unlock platooning benefits, but it poses significant challenges, including competitive fleet interests, data privacy concerns, high computational complexity, and the need for real-time decision making under uncertainties [S43].

Effective platoon coordination over large transportation networks requires optimization from various perspectives, including efficient route planning [S44], fair market incentives [S45], and departure time scheduling [S46]. These strategies maximize the number of platooning opportunities, enhancing the resulting benefits for fleet owners as well as society at large (Figure S13).

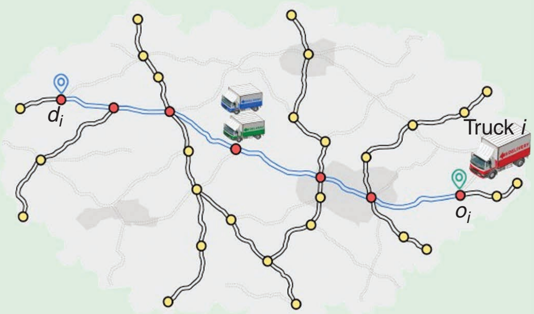

FIGURE S12 Hub-based platoon coordination. Trucks from different fleets with overlapping routes can be coordinated to form platoons at hubs.

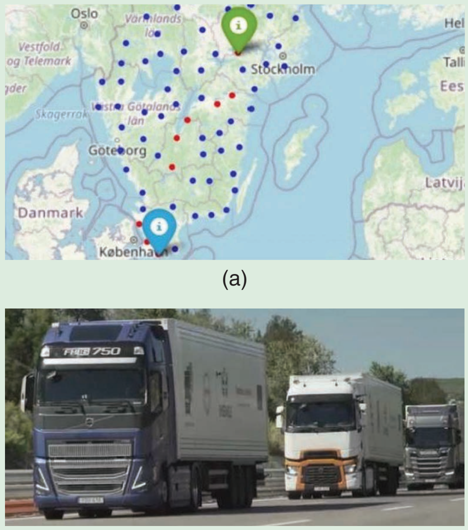

FIGURE S13 Large-scale platoon coordination on the southern Swedish road network. (a) The hubs along the route of a single truck are shown by red nodes, while other hubs are shown by blue nodes. An automated system suggests to truck drivers at which hubs it is most beneficial to form platoons with other trucks. (b) Trucks traveling in a three-vehicle platoon during an experimental evaluation of the system.

REFERENCES

[S43] T. Bai, A. Johansson, K. H. Johansson, and J. Mårtensson, "Large-scale multi-fleet platoon coordination: A dynamic programming approach," *IEEE Trans. Intell. Transp. Syst.*, vol. 24, no. 12, pp. 14,427–14,442, Dec. 2023, doi: 10.1109/TITS.2023.3298564.

[S44] S. H. van de Hoef, K. H. Johansson, and D. V. Dimarogonas, "Fuel-efficient en route formation of truck platoons," *IEEE Trans. Intell. Transp. Syst.*, vol. 19, no. 1, pp. 102–112, Jan. 2018, doi: 10.1109/TITS.2017.2700021.

[S45] F. Farokhi and K. H. Johansson, "A study of truck platooning incentives using a congestion game," *IEEE Trans. Intell. Transp. Syst.*, vol. 16, no. 2, pp. 581–595, Apr. 2015, doi: 10.1109/TITS.2014.2329317.

[S46] A. Johansson, T. Bai, K. H. Johansson, and J. Mårtensson, "Platoon cooperation across carriers: From system architecture to coordination," *IEEE Intell. Transp. Syst. Mag.*, vol. 15, no. 3, pp. 132–144, May/Jun. 2023, doi: 10.1109/MITS.2022.3219997.

for Research and Innovation, and the companies Scania and Ericsson.

ACKNOWLEDGMENT

This work was financially supported by Knut and Alice Wallenberg Foundation Wallenberg Scholar and Engineering the Interconnected Society Project grants; the Wallenberg AI, Autonomous Systems, and Software Program; Swedish Research Council Distinguished Professor Grant 2017-01078; Swedish Agency for Innovation Systems Grants AllDrive and Sweden4Platooning; the European Union Framework Program Horizon 2020 ENSEMBLE project and oCPS Innovation Training Network; and Swedish Strategic Research Foundation Grants SoPhy and SUCCESS.

AUTHOR INFORMATION

Karl H. Johansson (kallej@kth.se) is the Swedish Research Council Distinguished Professor in Electrical Engineering and Computer Science at KTH Royal Institute of Technology, 100 44 Stockholm, Sweden, and founding director of Digital Futures. He earned his M.Sc. degree in electrical engineering and Ph.D. degree in automatic control from Lund University. He has held visiting positions at the University of California, Berkeley; California Institute of Technology; Nanyang Technological University; Norwegian University of Science and Technology; and other institutions. His research interests focus on learning-based networked control systems and cyberphysical systems, with applications in

transportation, energy, and automation networks. For his scientific contributions, he has received numerous best paper awards and various other distinctions from IEEE, the International Federation of Automatic Control (IFAC), and other organizations. In addition to being awarded the title of distinguished professor by the Swedish Research Council, he has been named a Wallenberg scholar by the Knut and Alice Wallenberg Foundation and a future research leader by the Swedish Foundation for Strategic Research. He has also received the triennial IFAC Young Author Prize, IFAC Outstanding Service Award, and CSS Hendrik W. Bode Lecture Prize, and he has been an IEEE Control Systems Society (CSS) Distinguished Lecturer. His service to the academic community includes being president of the European Control Association; CSS vice president of diversity, outreach, and development; and a member of the CSS Board of Governors as well as of the IFAC Council. He has served on the editorial boards of *Automatica*, *IEEE Transactions on Automatic Control*, *IEEE Transactions on Control of Network Systems*, and many other journals. He has also been a member of the Swedish Scientific Council for Natural Sciences and Engineering Sciences. He is a Fellow of IEEE and the Royal Swedish Academy of Engineering Sciences.

REFERENCES

[1] "Energy technology perspectives 2020," Int. Energy Agency, Paris, France, Tech. Rep., 2020. [Online]. Available: https://www.iea.org/reports/energy-technology-perspectives-2020

[2] "Taxonomy and definitions for terms related to on-road motor vehicle automated driving systems," SAE Int., Warrendale, PA, USA, Tech. Rep., Jan. 2014. [Online]. Available: https://www.sae.org/standards/content/j3016_201401/

[3] F. Jiang, "Human-centric control design for safe and connected vehicles," Ph.D. dissertation, KTH Royal Inst. of Technol., Stockholm, Sweden, Apr. 2025. [Online]. Available: https://kth.diva-portal.org/smash/get/diva2:1942660/FULLTEXT01.pdf

[4] V. Narri, "Shared situational awareness under complex traffic scenarios," Licentiate thesis, KTH Royal Inst. of Technol., Stockholm, Sweden, 2025. [Online]. Available: https://diva-portal.org/smash/record.jsf?pid=diva2:1938150

[5] O. Grembek, A. A. Kurzhanskiy, A. Medury, P. Varaiya, and M. Yu, "Introducing an intelligent intersection," Univ. of California Inst. of Transp. Studies, ITS Rep. 13, 2018. [Online]. Available: https://escholarship.org/uc/item/2qm9h8jb

[6] T. Nyberg et al., "Share the unseen: Sequential reasoning about occlusions using vehicle-to-everything technology," *IEEE Trans. Control Syst. Technol.*, vol. 33, no. 4, pp. 1418–1431, Jul. 2025, doi: 10.1109/TCST.2024.3499832.

[7] D. Bertsekas and I. Rhodes, "On the minimax reachability of target sets and target tubes," *Automatica*, vol. 7, no. 2, pp. 233–247, 1971, doi: 10.1016/0005-1098(71)90066-5.

[8] M. Althoff, M. Koschi, and S. Manzinger, "Commonroad: Composable benchmarks for motion planning on roads," in *Proc. IEEE Intell. Vehicles Symp.*, 2017, pp. 719–726, doi: 10.1109/IVS.2017.7995802.

[9] P. F. Orzechowski, A. Meyer, and M. Lauer, "Tackling occlusions & limited sensor range with set-based safety verification," in *Proc. 21st Int. Conf. Intell. Transp. Syst. (ITSC)*, 2018, pp. 1729–1736, doi: 10.1109/ITSC.2018.8569332.

[10] B. Besselink et al., "Cyber–physical control of road freight transport," *Proc. IEEE*, vol. 104, no. 5, pp. 1128–1141, May 2016, doi: 10.1109/JPROC.2015.2511446.

[11] J. L. de Lagrange, "Mémoire sur la théorie du mouvement des fluids" *Nouveaux Mémoires de l'Académie royale des Sciences et Belles-Lettres de Berlin*, vol. 4, pp. 695–748, 1781.

[12] M. Papageorgiou, H. Hadj-Salem, and J. M. Blosseville, "ALINEA: A local feedback control law for on-ramp metering," *Transp. Res. Rec.*, vol. 1320, no. 1, pp. 58–67, 1991.

[13] A. Ferrara, S. Sacone, and S. Siri, *Freeway Traffic Modelling and Control*. Cham, Switzerland: Springer-Verlag, 2018.

[14] M. Čičić and K. H. Johansson, "Traffic regulation via individually controlled automated vehicles: A cell transmission model approach," in *Proc. Int. Conf. Intell. Transp. Syst.*, 2018, pp. 766–771, doi: 10.1109/ITSC.2018.8569960.

[15] R. E. Stern et al., "Dissipation of stop-and-go waves via control of autonomous vehicles: Field experiments," *Transp. Res. Part C*, vol. 89, pp. 205–221, Apr. 2018, doi: 10.1016/j.trc.2018.02.005.

[16] M. Barreau, M. Aguiar, J. Liu, and K. H. Johansson, "Physics-informed learning for identification and state reconstruction of traffic density," in *Proc. IEEE Conf. Decis. Control*, 2021, pp. 2653–2658, doi: 10.1109/CDC45484.2021.9683295.

[17] E. H. Lee and E. Lee, "Congestion boundary approach for phase transitions in traffic flow," *Transportmetrica B, Transp. Dyn.*, vol. 12, no. 1, 2024, Art. no. 2379377.

[18] M. J. Lighthill and G. B. Whitham, "On kinematic waves II. A theory of traffic flow on long crowded roads," *Proc. R. Soc. London, Ser. A. Math. Phys. Sci.*, vol. 229, no. 1178, pp. 317–345, 1955.

[19] P. I. Richards, "Shock waves on the highway," *Operations Res.*, vol. 4, no. 1, pp. 42–51, 1956, doi: 10.1287/opre.4.1.42.

[20] J. P. Lebacque, J. B. Lesort, and F. Giorgi, "Introducing buses into first-order macroscopic traffic flow models," *Transp. Res. Rec.*, vol. 1644, no. 1, pp. 70–79, 1998, doi: 10.3141/1644-08.

[21] M. L. D. Monache and P. Goatin, "A front tracking method for a strongly coupled PDE-ODE system with moving density constraints in traffic flow," *Discrete Continuous Dynamical Syst.*, vol. 7, no. 3, pp. 435–447, 2014.

[22] M. Čičić, "Modelling and Lagrangian control of mixed traffic: Platoon coordination, congestion dissipation and state reconstruction," Ph.D. thesis, KTH Royal Inst. of Technol., Stockholm, Sweden, 2021.

[23] M. Raissi, P. Perdikaris, and G. Karniadakis, "Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations," *Comput. Phys.*, vol. 378, pp. 686–707, Feb. 2019, doi: 10.1016/j.jcp.2018.10.045.

[24] M. Barreau, J. Liu, and K. H. Johansson, "Learning-based state reconstruction for a scalar hyperbolic PDE under noisy Lagrangian sensing," in *Proc. 3rd Conf. Learn. Dyn. Control*, 2021, pp. 34–46.

[25] M. Behrisch, L. Bieker, J. Erdmann, and D. Krajzewicz, "Sumo – Simulation of urban mobility; An overview," German Aerospace Centre, Cologne, Germany, Inst. for Transp. Res., Tech. Rep., 2011.

[26] L. Giorgi, "Prise en compte des transports en commune de surface dans la modélisation macroscopique de l'Écoulement du trafic," Ph.D. thesis, Institut National des Sciences Appliquées de Lyon, Villeurbanne, France, 2002.

[27] M. Čičić and K. H. Johansson, "Front-tracking transition system model for traffic state reconstruction, model learning, and control with application to stop-and-go wave dissipation," *Transp. Res. Part B*, vol. 166, pp. 212–236, Dec. 2022.

[28] T. Bai, Y. Li, K. H. Johansson, and J. Mårtensson, "Rollout-based charging strategy for electric trucks with hours-of-service regulations," *IEEE Control Syst. Lett.*, vol. 7, pp. 2167–2172, 2023, doi: 10.1109/LCSYS.2023.3285137.

[29] T. Bai, Y. Li, A. A. Malikopoulos, K. Henrik Johansson, and J. Mårtensson, "Distributed charging coordination for electric trucks under limited facilities and travel uncertainties," *IEEE Trans. Intell. Transp. Syst.*, vol. 26, no. 7, pp. 10,278–10,294, Jul. 2025, doi: 10.1109/TITS.2025.3550035.

[30] D. P. Bertsekas, *Lessons from AlphaZero for Optimal, Model Predictive, and Adaptive Control*. Belmont, MA, USA: Athena Scientific, 2024.