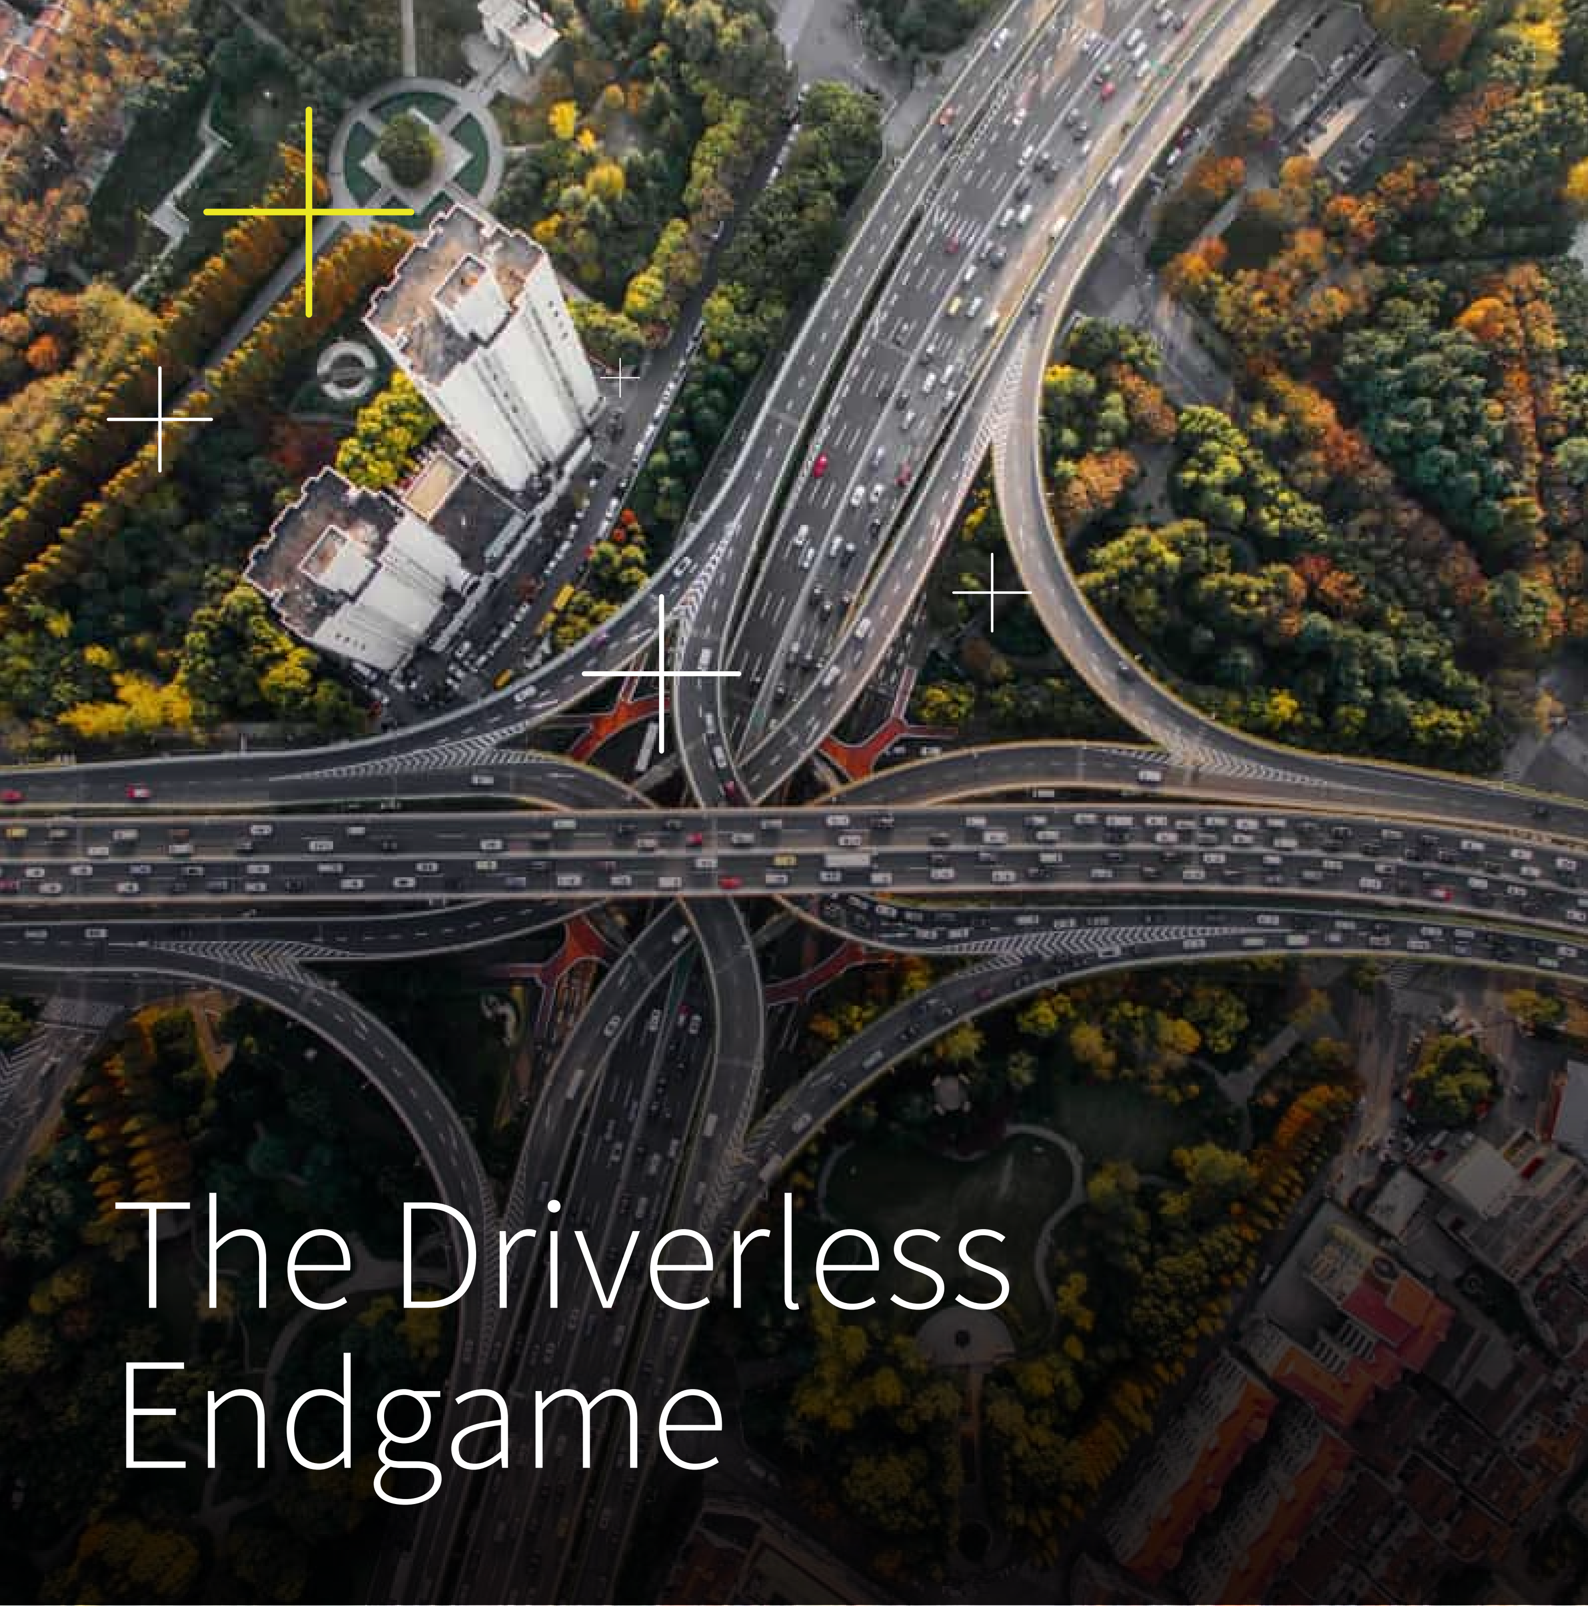

The Driverless Endgame

Policy and Regulation for Automated Driving

HATCH

Key Acronyms

ADAS

Advanced Driver Assistance System

A system that assists a human driver to perform the entire dynamic driving task (DDT) but requires active human oversight and presence.

ADS

Automated Driving System

A system, combining hardware and software, collectively capable of performing the DDT on a sustained basis, regardless of whether it is limited to a specific Operational Design Domain (ODD).

ADS-H

ADS-Human

A system that permits a human fallback driver to disengage the ADS and assume driving the vehicle when the vehicle is in its ODD.

ADS-V

ADS-Vehicle

A system that does not permit a human fallback driver to disengage the ADS and operate the vehicle when the vehicle is in its ODD.

ADS-DV

ADS-Dedicated Vehicle

An ADS-equipped vehicle designed for driverless operation during all trips within its given ODD

The Glossary in Chapter 17 has a complete list of acronyms used in this white paper.

Two Important Disclaimers

Almost all the recommendations herein align with emerging global best practice, as we understand it ("emerging" because actual deployment experience is scant or non-existent). A few, however, are novel; see, for instance, our recommendations about the introduction of the terms ADS-H and ADS-V in Section 4.1.3. In all cases, these recommendations depend on preliminary analysis of evolving technology, evidence, and approaches, and in most cases no consensus has yet emerged. All proposals we have made deserve to be tested by stakeholders and experts before any regulator proceeds with them.

Similarly, this white paper offers insights into policy to inform consideration of a future legislative and regulatory framework for driving automation. While the paper provides several directional recommendations regarding changes in policy and regulations, it does not provide legal advice, as the authors are not qualified to offer such. All recommendations herein should be independently reviewed by appropriate experts, legal and otherwise.

List of Recommendations to the Regulator

Language for Automated Driving	**1.**	In a regulatory context, use the terms ADAS, ADS, ADS-H, ADS-V, and ADS-DV
	2.	In a regulatory context, avoid use of SAE3016 engineering level classifications, e.g., "Level 4 vehicle"
Liability, Responsibility, and Insurance	**3.**	Establish a clear, unambiguous regulatory classification standard for driving automation that recognizes Driver fallback and System fallback, exactly as described by SAE3016
	4.	Prohibit any ADS that can be enabled or that can remain enabled outside of its ODD
	5.	Require generous and specified lead-times when the ADS anticipates the end of an ODD
	6.	Divide the definition of ADS into three categories, namely ADS-H, ADS-V, and ADS-DV
	7.	Audit the jurisdiction's roads to establish which ones are appropriate for ADS operation, and do so in advance of licensing any ADS-enabled vehicles
	8.	Establish a clear event-line of responsibility for vehicles equipped with ADS, such that it is unambiguous who is responsible for the vehicle at all times
	9.	Bias liability toward the owner/operator of ADS-equipped vehicles, to the degree possible
	10.	Consider how to prevent the marketing of an ADAS or ADV vehicle with language that misleads regarding the extent of its capabilities
	11.	Require OEMs to execute a contract that outlines the necessary and sufficient conditions for registration, and require evidence of sufficient insurance
Digitalizing the Rules of the Road	**12.**	Refine the rules of the road such that they are identical for humans and ADS
	13.	Develop a master source for the rules of the road in digital form
	14.	Consider the intention and direction of the ISO METR project
Vehicle User Requirements	**15.**	Identify, as part of any vehicle-user requirement, the entity responsible for meeting that requirement, whether human or corporation, and prohibit the operation of vehicles without an onboard responsible human agent until this identification is complete
	16.	Define a minimum expected self-checking test regime
	17.	Define the regulatory regime for teleoperation of ADS-equipped vehicles
Training and Licensing	**18.**	Amend the driver-licensing process to include explicit discussion of ADAS (and, in time, ADS), and to emphasize that it is the driver's responsibility to understand any ADAS or ADS features they employ while operating a motor vehicle
	19.	Begin to design a consistent licensing and registration process specifically for ADS

Enforcement and Emergency Response

20. Require that:

a. ADS-DV vehicles feature a unique external marking to facilitate remote identification
b. Enforcement officers may, upon identification, issue electronic citations to ADS-DV vehicles without apprehending the vehicle
c. ADS-DV vehicles feature a teleoperator that enforcement and emergency personnel may contact
d. Teleoperators must comply with enforcement officers
e. ADS-DV and ADS-V vehicles feature an override system by which they may be directed to safely come to a stop and disable themselves

21. Adopt law-enforcement interaction protocols specifically for ADS vehicles that specify how to:

a. Recognize a vehicle operating in ADS mode
b. Safely disengage it and know that it has been disengaged
c. Communicate with a teleoperator
d. Make these protocols available via the Internet to facilitate fast and simple access by law enforcement

22. Harmonize law enforcement interaction protocols across jurisdictional boundaries

23. Determine rules and procedures for responding to an emergency affecting ADS-DV vehicles where physical presence of an emergency crew at the site of the emergency is required

Aftermarket Modifications

24. Restrict aftermarket modification of an ADS to be made only by the entity or entities that is/are held liable for the vehicle when its ADS is engaged

25. Prohibit any ADS modifications that original equipment manufacturers request to be so prohibited, subject to a test for reasonableness

Automated Trucking and Cooperative Truck Platooning

26. Prioritize the regulation of ADS to permit its use in trucking and cooperative platooning

27. Permit deployment of ADS for trucking and cooperative platooning in well-understood and tightly-controlled stages

28. Insist upon appropriate measures to ensure timely and consistent human response to truck ADS disengagement, and verify that they are respected

29. Adopt a conservative program of registrations for ADS deployment, deploying only as fast as fleets can unequivocally prove themselves to be operationally safe at each stage

30. Acknowledge that it is inappropriate and impractical to prove road safety for ADS via road-testing alone

31. Retain existing rules for hours of service until critical performance milestones are met

32. Prohibit truck ADS-DV deployment until a pre-determined proof of safety program for ADS-H or ADS-V driving has completed

33. Permit platooning changes—longer trains, mixed vehicles, removal of follower drivers, and so forth—in measured steps with strict safety tests

34. Harmonize platooning rules across jurisdictional boundaries to the extent possible

35. Limit platoons initially to two vehicles, sharing a common registration, and common ADS

Harmonization	**36.** Do not legislate OTA updates as the default for cross-border changes to rules of the road **37.** Pursue the highest possible degree of harmonization of driving-automation regulation, for all relevant parties within contiguous regions
Data Collection	**38.** Determine which ADS-related operating data to collect, as well as the data's format, sampling frequency, availability, permitted uses, and schedules for retention and destruction **39.** Consult with data and privacy experts in determining which ADS-related operating data to collect, and under what circumstances to make it available, and to which parties **40.** Determine a way to ensure sensitive collection, and use, of violation data
Loading, Unloading, Curb Space, and Parking	**41.** Provide guidance to help cities understand and adjust to required changes in parking and loading matters
Public Mobile Robots	**42.** Consult with accessibility, logistics, municipal, planning, and robotics experts to design a model code

Executive Summary

Driving automation is a technology that assists with the driving task, up to and including performing it entirely. Popularly and misleadingly known as "autonomous vehicles", this technology will alter, reduce, and eventually eliminate the role of the human driver. In jurisdictions around the world, laws and regulations for the use and operation of road vehicles will require revision to ensure the safe and timely deployment of this technology.

This white paper aims to:

1. Assess legislative and regulatory barriers and gaps that may impact prompt, safe deployment of driving automation on public roadways
2. Assess options for changes to the transportation regulatory environment to address these barriers and gaps
3. Support this technology's potential to increase the safe, efficient movement of people and goods
4. Offer recommendations to the regulator on how to achieve these goals

These assessments and recommendations rely in part on research and analysis—specifically, literature reviews, jurisdictional scans, and interviews of topic experts—and in part on the authors' experience as specialists in both policy analysis and driving automation.

This white paper comprises reviews and recommendations organized into several policy matters, as follows.

A. Language for Automated Driving

Discussion of automated driving, in both specialist and popular venues, relies on language taken from the Society of Automotive Engineers (SAE)'s taxonomy of driving-automation systems (SAE 3016_202104). This paper proposes that this language is useful for engineers but unhelpful for regulators and proposes a very strict usage of this SAE terminology for discussing vehicles that feature automated-driving capabilities. The paper argues that careful usage helps by providing for unambiguous assignment of liability to a human driver or to the ADS, as appropriate. These terms avoid language that the SAE itself has recommended not be used.

B. Liability, Responsibility, and Insurance

There are fundamentally two forms of driving automation. One aims to assist a human driver by automating elements of the dynamic driving task, known as Advanced Driver Assistance Systems (ADAS). The other aims to carry out the entire dynamic driving task itself, with possibly very occasional and preferably *absolutely no* resort to a human driver as fallback, known as Automated Driving Systems (ADS). There are, or are expected to be, multiple, competing instances of each.

To remove ambiguity regarding responsibility and liability for the dynamic driving task, this paper draws a single, distinct line between when an ADS is engaged and when it is not. A vehicle with engaged ADS is one case; all other driving situations, including engagement of ADAS, form a separate case. This dividing line is based firmly on advice from the SAE3016 standard that guides the formal language for this technology. Excepting product failure, infrastructure failure, or *force majeure*, this paper follows the implication of the SAE standard to argue that the ADS provider is fully responsible when an ADS is engaged, and the human driver is fully responsible when an ADS is not engaged (engagement of ADAS is not an instance of ADS). This paper argues that complexities that may arise, due to exceptions or the fact that ADS provision involves multiple parties, is a matter for insurance subrogation rather than government regulation. The paper also argues for strict rules regarding the engagement of ADS and for marketing promises that are not backed by the acceptance of liability.

The benefits of this approach include the minimization of ambiguity regarding crashes, violations, and settlements; relegation of the complexity of product-related financial settlements to insurance subrogation; and consignment of liability for the dynamic driving task to the human driver until such time that the providers of ADS are prepared to accept full liability for their products. An additional benefit of this approach is that it will permit the policy authority to generate a body of regulations that does not depend upon complex language, nor on the nuances of driving automation system designs.

C. Digitalizing the Rules of the Road

The rules that a human driver is expected to follow while piloting a motor vehicle must also be embedded into the software of an ADS, such that a vehicle with its ADS engaged would perforce comply with those rules. Currently, the rules of the road as expressed in regulation often rely on language that refers to human cognitive, social, or even emotional capabilities, such as *attention, due care, prudent*, or *reasonable consideration*. Since all such rules need to apply to both human drivers and ADS, they must be rewritten in a way that is neutral to the nature of the driver, but still consistently enforceable on the road and consistently interpretable in a court of law.

Such a reformulation of the rules of the road would have to be encodable in a software system and be interpretable for a court that might be considering both human and machine behaviour in the same case. Our recommendation is to write these rules as machine code first and only then to interpret them for human instruction, for enforcement of human and machine drivers, and for the courts. Such a rendering of a digital code for human use might appear quite similar in language to the current analog code, but enforcement would rely on the base digital construction. In other words, rather than proposing that existing regulations be digitally twinned, this paper instead proposes multiple analog twins to digital, harmonized road-use legislation. In so proposing, this paper aligns with an emerging best practice, namely METR, a proposed ISO standard for managing digital traffic rules.

The first benefit of this approach is to have a single, verifiable, translatable body of rules that are prepared for an increasingly digitalized and connected traffic system, one that could not only be provided to ADS-equipped vehicles, but also to connected ADAS-equipped vehicles under the supervision of a human driver.

The second benefit is to make more tractable the harmonization, encoding, updating, and enforcement of rules of the road that apply to both human and machine drivers.

The third, and likely most critical, benefit is that as driving becomes increasingly digitalized, what is now an entirely analog world of rules and enforcement will transform into an entirely digital world of rules and enforcement. The sooner governments make the switch to a digital master from an analog master, the safer and easier this transformation will be.

D. Vehicle User Requirements

This paper argues that in all cases of vehicle operation there must always be a responsible human operator. That is usually a driver in the driver's seat of a vehicle, whether as the active driver or as a fallback driver. The important exception is a class of ADS called ADS-DV (Dedicated Vehicle, i.e., a vehicle in which the automated driving system performs the entire dynamic driving task, corresponding to all SAE Level 5 vehicles and some Level 4s). In this latter case, in which a moving vehicle has no responsible human driver inside, such a vehicle must have an increased ability to self-check its components for functional operation, especially its sensor array, and there must always be an available teleoperator. Regardless of the level and type of attention such a teleoperator provides, it is critical that the teleoperator always be able to address vehicle user requirements. The regulator must set new rules regarding vehicle user requirements, as such requirements cannot be abrogated.

The first benefit of this approach is to make it unlikely for a vehicle to be in motion on public roadways without provision for road safety due to shortfalls in vehicle preparation or systems readiness, load securement, or passenger securement. The second benefit is to ensure that ADS-DV vehicle designers, fleet operators, teleoperators, and insurers consider very carefully their contribution to, and liability for, vehicle systems and safety beyond following the rules of the road.

E. Training and Licensing

As elements of the dynamic driving task are increasingly undertaken by ADAS, the nature of driving is changing from direct operation to oversight. Because this change has heretofore been incremental and slow, most drivers today have adjusted to be sufficiently safe, even without formal training. Nonetheless, human driver behaviour while using ADAS-equipped vehicles often exhibits faults and unintended consequences related to attention, overconfidence, overreliance, misuse, unawareness, and selective disconnection of warning components. Hence, this paper recommends regulation to require that sellers, leasers, and renters of ADAS- and ADS-equipped vehicles offer education in the engagement, disengagement, and use of ADAS and ADS features to the buyers and renters of vehicles so equipped.

The paper recommends that regulators amend their driver licensing processes to include explicit instruction regarding ADAS—and, in time, ADS—and that this instruction emphasize that it is the driver's responsibility to understand not only any ADAS or ADS features they engage while driving, but also that they are liable for misuse of any driving technology or incorrect engagement of, or reliance on, ADS outside of its intended operational design domain.

At present, training and licensing regimes are for human drivers only. The implication of pending technology is that such regimes will become insufficient, because the act of driving a car (with or without ADAS) is very different from the responsibility of supervising a robot (ADS). Accordingly, this paper recommends that the regulator begin to understand and design a licensing and registration process specifically for ADS. Such a process should include both private vehicle owners and fleet owner/operators of service vehicles.

The benefit to this approach is to train new drivers for the new increasingly automated vehicle fleet of the future, bring incremental awareness to existing drivers as they encounter partially automated vehicles, and to prepare for the pending arrival of ADS fleet operators.

F. Enforcement and Emergency Response

This paper expects that both enforcement and emergency response will become more complex in the coming decades, when non-automated, partially automated, and fully automated systems share the same roads. There will be unintended interactions between vehicles with and without human drivers, new vehicle designs, more sophisticated software, new vehicle materials and energy systems, and potentially novel forms of illegal vehicle use. While one might expect the demand for enforcement and emergency responses to gradually decline, such an outcome will take many decades to materialize. Even if demand declines, its complexity is unlikely to diminish.

The process of detecting, identifying, and citing a motor vehicle for a moving violation will increasingly depend on remote methods rather than physical apprehension of the vehicle. Still, there will be cases where it will be necessary to apprehend an ADS-equipped vehicle for a violation (whether vehicle, moving, or criminal). This paper recommends that, to ensure that traffic enforcement personnel be granted the ability to safely force such a vehicle to stop, pull over, and disable its ability to move, such vehicles be required to feature a teleoperator with whom enforcement personnel may directly communicate. The paper also discusses interactions between enforcement officers and passengers in a circumstance where there is no “responsible driver” among the vehicle’s passengers. These interactions will require new training for enforcement officers and new considerations of privacy expectation for passengers.

Emergency responders must be able to pre-identify attributes about a vehicle prior to attempting entry, apprehension, or extinguishing a fire. This paper describes ways to facilitate this pre-identification, via real-time identification schemes and direct connection to teleoperators.

The benefit of implementing these recommendations is that, absent such, it will be impossible to manage automated driving systems given current analog enforcement and rescue methodologies.

G. Aftermarket Modifications

Driving automation relies on sensors, effectors, and software. These elements, taken together, comprise a system of mutual reliance. Sensors need to be aligned, effectors need to remain connected and responsive to the software, software needs to be matched with the sensors and effectors for which it was prepared, and the software must remain up to date. Any alteration to these parts, or to the integration among the components risks making a system—especially an ADS—unreliable.

This paper recommends that any ADS modification, including modifications to its sensors, effectors, and software updates be made only by, or under supervision of, that entity that is held liable for the vehicle when its ADS is engaged, and that this prohibition be strictly enforced. The paper acknowledges that this recommendation is in tension with the “right to repair” and explores some approaches to reducing that tension.

H. Automated Trucking and Cooperative Truck Platooning

This paper argues for ADS-V trucks with an attentive fallback driver as a route to safer truck traffic, better working conditions for drivers, and eventually extended hours of service (depending on the degree of fallback attention required). The degree of these benefits would depend on the extent of the ADS ODD, the reliability of re-capturing the fallback driver’s attention, and the capability of bringing the vehicle to a safe-stop condition in the counter-event of a non-attentive driver.

To achieve this end, the paper recommends a program of registering ADS-equipped trucks in well-understood and tightly-regulated stages—encompassing geographies, fleet size, ODD definitions—such that the advancement to a new stage relies on pre-determined safety conditions during a previous stage. The paper does not recommend, or even contemplate, HGVs without a human fallback driver (ADS-DV) for the foreseeable future, on expectation that a full end-to-end ODD is very difficult and unlikely to emerge in the next few decades (except perhaps for highly-constrained hub-to-hub transportation), and consequently see less need to regulate for HGV ADS-DV at this time.

Cooperative truck platooning conceptually starts with a lead truck and a follower truck or trucks, each with an attentive driver, such that the ADAS on follower trucks locks into the lateral and longitudinal acceleration systems on the lead truck and follows very closely. Follower drivers would not need to steer or brake, but would need to remain attentive to take over the dynamic driving task at short notice. Today, this simple level of short-train platooning has many constraints but, within those constraints, shows promise. This paper recommends that platooning be constrained to two trucks and at least one driver with tested and approved ADAS for the foreseeable future.

It is also the case, conceptually, that platooning could extend to multiple vehicles from multiple companies and multiple jurisdictions, with heterogeneous platooning software and communication systems and even with some vehicles in a train without a driver in the cab. The paper asserts that such hybrid systems are very far in the future and should not absorb regulators’ attention until short-train platooning enjoys a level of safety approaching that of commercial airliners; a standard that will be hard to match.

I. Harmonization

This paper argues that harmonization of policies and regulations relative to roadway vehicle movement and operation is a critical step for reasons of safety, enforceability, innovation, inspection, and commerce, and firmly recommends harmonization to every degree possible across all subnational jurisdictions for a uniform system within a nation; and, beyond these, urges a significant effort be made to harmonize across contiguous nations.

Harmonization over such a broad spectrum of jurisdictions and crossing national boundaries is a complex exercise that may take years, but nonetheless the paper argues that regulators should seek this outcome, regardless of how long it may take to realize.

J. Data Collection

Data collection, application, and privacy are controversial matters, perhaps especially so for the systems that support driving automation. This controversy will be especially fraught during the long transition to ADS-dominant roadways because records of how, why, where, and when passengers and goods move holds virtually boundless personal, commercial, management, and investment value. This paper recommends that regulators consult with privacy experts to determine which ADS-related operating data to collect, as well as the data's format, sampling frequency, availability, permitted uses, and schedules for retention and destruction. This work should begin early, in preparation for the long period of social-political debate that will ensue.

K. Loading, Unloading, Curb Space, and Parking

There is already ample demand for vehicle storage space near key trip generators, and for pick-up-and-drop-off space at them. As vehicle automation spreads, this demand will increase significantly, as accessing these spaces may no longer require human supervision. Additionally, there is potential for disorder, as the human proclivity to stop at any curbside as desired could stress automated systems that cannot negotiate such stoppage easily. There is related ISO standards work that will help to address these issues; the paper suggests that regulators provide guidance, stemming from those standards, to their municipalities, on how to prepare their infrastructure and orchestration methods to accommodate these future pressures.

L. Sidewalk Robots

In this chapter we discuss public mobile robots (PMRs), which share some aspects and technologies with driving automation for road-worthy motor vehicles. PMRs possess many unique features which, in this paper's view, require very different governance. Colloquially known as sidewalk robots, PMRs have begun to operate in maintenance and delivery services in communities around the world. There are now likely more than 100 makers of PMRs for applications on urban footways. Most are intended for last-mile delivery, but some plow snow, remove litter, or act as mobile surveillance devices.

The capability of these machines to traverse walkways, bikeways, and road shoulders while being operated via a combination of ADS and (human) radio teleoperation creates a new traffic circumstance. Pedestrian footways and trails are highly variable, often severely constrained in width, and have principally been used by humans with a wide variety of familiarity or comfort with PMRs. Determining the appropriate balance of uses is a complicated regulatory matter.

A key concern that the paper outlines is the matter of managing traffic from multiple operators with fleets of numerous robots each independently tasked and directed among non-involved pedestrians and cyclists in public spaces. This environment of deploying PMRs is very different from their deployment in factory or agricultural workspaces wherein the humans involved are trained to collaborate with the automated mobile robots (AMRs).

This paper describes a draft ISO standard intended to address public mobile robot behaviour, certification, and traffic orchestration for safe and workable operation in public spaces and recommends that regulators give it due consideration.

Table of Contents

List of Tables and Figures

About the Authors

Andrew Miller

Co-Author

Andrew Miller is a city-builder, urbanist and futurist specializing in mobility and transportation planning, land use, and stakeholder relationship building. He has a twenty-year career that spans government at the provincial and municipal levels, as well as private-sector advisory. As a Senior Leader and Consultant at Hatch, Andrew delivers a variety of projects, with emphasis on urban mobility and innovation, including masterplanning mobility networks for new neighbourhoods; transit-oriented development; and future-proofing urban designs to consider innovative transport technology.

Prior to joining Hatch, Andrew served as the Associate Director, Mobility for **Alphabet's Sidewalk Labs**, where he served as Toronto mobility lead for the firm's flagship Quayside smart-city project, a program that integrated people-centered urban design with cutting-edge mobility technology. In that role, he helped to develop the mobility program for this proposed "city of the future", including the design of pick-up-and-drop-off infrastructure for passenger-focused "robotaxis" and loading docks for autonomous freight vehicles, and in-depth evaluations of automated-shuttle providers for licensing, partnership, or purchase by Sidewalk Labs. Additionally, in 2018 and 2019 Andrew served as an advisor to the **Canadian Senate's Vehicle of the Future Advisory Group**.

He may be contacted at andrew.miller@hatch.com

Bern Grush

Co-Author

Bern Grush is a Canadian systems-design engineer with a reputation for excellence in innovation, research, writing, teaching, and speaking in the area of new mobility, especially regarding automation and how this automation might be regulated, governed, socialized, and deployed.

Bern is the lead author of the 2018 textbook *The End of Driving: Transportation Systems and Public Policy Planning for Autonomous Vehicles*, a textbook on the role, opportunities, and implications of driverless vehicles for cities, which Alain Kornhauser, the Princeton University Faculty Chair in autonomous vehicle engineering, described as "a serious publication aimed at researchers, planners, policymakers, and practitioners. Most highly recommended." This book was the product of two years of research on matters of governing, socializing, and deploying large, public projects of autonomous passenger vehicles, making Bern one of the **world's foremost experts on the public-policy implications of automated driving**.

Since 2019 Bern has been the project leader for the drafting of **ISO 4448 public mobile robots and automated pathway devices**. This draft ISO standard sets out the required data and procedures for managing specific aspects of the operation of fleets of automated passenger and goods vehicles on roadways, as well as the operation of disparate fleets of automated mobile service robots (e.g., delivery, snow plough, maintenance, or surveillance robots) in sidewalks, bicycle lanes, road shoulders, and crosswalks.

He may be contacted at bern@urbanroboticsfoundation.org

Victoria Webster

Researcher

Victoria is a civil engineer and an urban planner, specializing in transportation planning and economics. As a Consultant and Project Manager at Hatch, Victoria works with public and private sector clients to deliver a range of projects pertaining to transit-oriented development strategy, land value capture, station area design, transit masterplanning, and design future-proofing.

Her technical background drives a quantitative evidence-based approach she uses to solve complex problems. While at Hatch, Victoria has developed successful mobility and development strategies for major projects and clients across North America.

Erin da Silva

Researcher

Erin is a civil engineer and urban planner with experience in road and stormwater design and inspection, transport policy and socio-economic analysis, active mobility planning, and transit-oriented and neighbourhood development.

She works within the confluence of civil engineering and urban planning, where she provides strategic insight and consultation services. She has worked on projects involving the improvement of freight corridors in South Africa, as well as the Park Lawn GO Transit Hub, the City of Calgary Green Line LRT, and the Gautrain Capacity and Alternate Revenue Study.

Hatch

Hatch is a Canadian-based, employee-owned, multidisciplinary professional services firm that delivers a comprehensive array of technical and strategic services, including engineering, consulting, project and construction management, information technology, and process development. We have served clients for over six decades with corporate roots extending over 100 years and project experience in more than 150 countries around the world. With over 9,000 people in over 65 offices, we have more than $35 billion in projects currently under management. Our constant aim is to deliver unprecedented business results for our clients.

Within our Urban Solutions group, we provide technical and strategic consulting services to support all levels of government and the private sector. We are leading advisors on a variety of issues relating to the intersection of mobility, infrastructure, and urban living, with particular expertise in transit-oriented development; economic and financial analysis; place-making and urban development; and technological innovation.

For more, please visit www.hatch.com

The Urban Robotics Foundation

The Urban Robotics Foundation (URF) is a non-profit Canadian foundation established to help cities anticipate and prepare for the deployment of automated fleets of passenger and goods vehicles as well and public mobile robots. URF's work includes systems of design, thought and education; examples include software for managing micro-incentives to increase transit ridership; the textbook The End of Driving; an EU MOOC to promote zero-car ownership communities; and the drafting of the ISO 4448 standard for public mobile robots. URF focuses on the intersection of Human Factors Psychology and Systems Design Engineering to create solutions for liveable cities.

For more, please visit www.urbanroboticsfoundation.org

Preface

Many unreliable studies and forecasts have been written to predict how the automated vehicle will work, how it will be used, and how it will address a multitude of transportation problems. Less has been written about how automated driving might change our cities or how we might prepare cities for their arrival. Smaller still is the literature regarding how we might regulate the use of these vehicles. In this slim text, we hope to contribute a usable outline for this latter issue.

In Canada, the jurisdiction we are most familiar with, and in many other places around the world, there is a growing gap between technology and its governance, particularly regarding artificial intelligence (AI). This may be because of the relative novelty of AI technologies, or perhaps because hype-filled social media makes us more of aware of this gap. Consider the relatively small size of the audiences of tech-imaginative publications such as *Popular Mechanics* in the 1960s–1980s compared to the audience of Twitter, Facebook, and LinkedIn today.

In any case, the gap between automated driving technology and its governance will be one of the most difficult problems in this domain to solve. After all, automobility pervades almost everything—how we live, where we live, how we design cities, how we consume, our physical health, the environment, how we project ourselves in society, and much more. This means that the built form evolved for the automobile and the installed fleet of human-driven vehicles will not only take decades, perhaps a century, to transform, but the two systems—human drivers and software drivers—will necessarily have to co-exist for the duration of this change. In proposed cities such as Neom or Woven, the problem will be trivial compared to the change necessary in already-existing cities like Riyadh or Tokyo.

Difficult as the problem is, the flood of hype and polarized exaggeration makes it much worse. On one day, a report in the media tells us that fully automated vehicles are coming in two years. The next day, we hear that fully automated vehicles are not actually possible. The day after, we are told that all that was wrong, and there will be fully automated trucks this year after all.

Any sensible person knows to consider predictions carefully, but the manner and number of predictions; the breathless accounts of the value automated driving will unlock; and the ongoing revisions of claims regarding dangers, pitfalls, speed-to-realization, and unintended outcomes heighten the crisis for governments. Whether subsidizing EV purchases, installing charging networks, or widening roads, decision-makers are undeniably confused and likely frustrated—or even frightened—by this confusion.

These decision-makers are not at fault. Governments are always constrained by past decisions and the legislative structure which permits them to govern. By charter or statute, some governments cannot make rules about vehicles under a certain size, weight, or speed. Others cannot address vehicles that lack an on-board human operator. Changes in governance are needed before such legislators can even legislate.

Considering the full array of promised technical innovations and disruptions, there are uncertainties in the likelihood and ranges of arrival times for automated driving. As a base case, consider that we have stable estimates for the arrival of electric vehicles at scale; we know roughly how fast the EV fleet can be expected to grow and how charging infrastructure for that fleet will need to be distributed. However, because we are less sure of battery charging speeds in a few years, we do not know how to size these charging distribution nodes, so we need to find solutions for expandable nodes in a way that the real estate involved can be temporarily repurposed and the power servicing those nodes is easily expanded. There are several solutions to this problem, which have been enumerated, and described in detail; they will need to be regulated, planned, and then rolled out as necessary.

Our understanding of automated driving is nowhere near as advanced. We do not have reliable estimates for the arrival process of vehicles with automated driving systems. We are unsure about the spatial or weather-condition extent of the ODDs (operational design domains) for those vehicles, the demand curve for using them for mobility-on-demand (robotaxis), and the demand curve for owning such vehicles. Nor do we understand how they will mix at scale with traditional vehicles, i.e., those that still require a licensed human driver to operate them.

And that means we also do not understand the requirements for such matters as training, licensing, enforcement, and liability.

We are just now getting the first hint of how little we know; at time of writing, 50 robotaxis in San Francisco are making enough errors to cause sufficient traffic delays for the continuation of the pilot program that permits them to operate to be questioned. While we are confident the pilot will continue, whether in San Francisco or elsewhere, this indicates how difficult this problem is. AVs have moved from being *almost ready* from a technical perspective to being *hardly ready* from a regulatory perspective.

For that reason, we offer this white paper. It is our hope that by thinking through the problems that an attentive regulator faces, or will face shortly, we can help to narrow this gap. This matters because this technology, whether it arrives fast or slow, will bring enormous consequences. We hope to maximize the value of these consequences and minimize the ones that are unintended.

1

Approach

1.1 Structure

After this initial chapter on our approach, the study begins in earnest in Chapter 2 with an overview of the present moment in the development of automated driving. Using the language of Gartner's "Hype Cycle", we argue the field is presently in the so-called Trough of Disillusionment, where received opinion has sensibly rejected the overstated claims that have been so prevalent, but in the process has indulged heavily in suspicion of this technology, its promise, and its development cycle. We believe suspicion will rightfully hang over this technology for the next few years. We also believe that this skepticism will ultimately moderate, even though the technology will become widespread without ever reaching the levels described with such breathless excitement in 2010–2020.

Having provided this understanding of the state of the field, we begin our recommendations in Chapter 3 with a discussion of the language and terms that we will use throughout our recommendations to describe vehicles with automated-driving capability. We find that most language used by others, while inspired by the Society of Automotive Engineers' (SAE) influential work on the subject, is not only used in ways that the SAE clearly warned against, but also is unhelpful for regulators[1]. We provide a closely-defined overview of the SAE terminology, eschewing careless use of the sort frequently found in the press and some literature, with careful elucidation where we deem it necessary. We argue that such usage helps by providing for unambiguous assignment of liability to a human driver or to the automated system, as appropriate. These terms avoid language that the SAE itself has recommended not be used.

We then consider what we think are the key policy areas that regulators must address. We devote a chapter to each issue, making up chapters 4 through 14. Most of these chapters share a similar structure: they begin by outlining the matters at issue within that domain that we believe require consideration and are followed by our policy analysis and recommendations.

The chapters are as follows:

Responsibility, Liability, and Insurance

The advent of automated-driving systems has implications for how we assess responsibility and liability for incidents involving vehicles equipped with such systems. We consider how regulators should govern the insurance sector as it allocates liability among the various parties that might conceivably bear it: the driver, the owner, the operator, and the manufacturer of the vehicle and/or its component parts and systems including software, map, and communication subsystems.

Vehicle User Requirements

Prior to the emergence of automated driving, the driver of a vehicle was assumed to be the party responsible for the vehicle's use and was required to oversee all aspects of its operation. Examples of such areas of oversight include the behaviour of other occupants, the safe stowage of goods aboard, and the functionality of vehicle parts, such as mirrors, lamps, and signals. We consider how these allocations of responsibility might change in vehicles equipped with ADS. Additionally, we discuss the appropriateness of mandating the use of particular ADAS technologies that would render a vehicle inoperable if the ADAS determines the driver is incapable of operating the vehicle safely.

1 SAE 3016-202104

Rules of the Road

Today, in most places around the world, vehicle traffic legislation does not contemplate the dynamic driving task being undertaken by any entity other than a human, because this legislation predates the advent of automated-driving technology. We suggest possible changes to regulations to acknowledge and provide rules for drivers whose vehicles are equipped with ADAS, to take advantage of the possibilities for efficient-yet-safe operation.

Training and Licensing

ADS-DV vehicles, if and when they arrive, will make no demands of a human operator, and their occupants will not need to be licensed to drive, any more than passengers in taxicabs today do. Between that day and this, though, we will have vehicles equipped with ADAS, ADS-H or ADS-V, which offer drivers new capabilities that make the dynamic driving task easier. Some ADAS, at the time of writing, are already sufficiently sophisticated to maintain a car on a highway at a constant speed in a single lane, and under ideal conditions require only human oversight, attention, and hands-on-wheel, but not human operation. However, these systems fall far short of ADS capability, and their misuse can, and has already, led to incidents and avoidable deaths. We investigate whether ADAS and ADS like these require human operators to be trained and licensed above and beyond what regulators typically require today.

Enforcement and Emergency Response

Driving automation has the potential to complicate the enforcement of appropriate driver behaviour, the investigation of road incidents, and intervention in cases of emergency. We discuss possible regulatory changes to give law enforcement and first responders the appropriate tools to keep these tasks from becoming more burdensome.

Aftermarket Modifications

We foresee a tension between the right of an owner to modify, improve, or repair their possessions, the right of other road users to have confidence that other vehicles on the road will operate in a safe and consistent manner, and the desire of manufacturers—who will be judged by the behaviour of their products—to ensure their sophisticated ADAS- and ADS-equipped vehicles operate as intended. We explore this tension and recommend possible policy responses to it.

Automated Trucking and Cooperative Truck Platooning

It seems likely that ADS will arrive first in the trucking sector, given the predictability of highway driving versus driving in other environments and the ongoing shortage of drivers in most jurisdictions relative to demand. We consider how the current regulatory, policy, and enforcement regime of goods movement should, or should not, change to take driving automation into account.

Harmonization

Jurisdictions have a responsibility to regulate in a fashion that is responsive to their communities' own views of what will promote the general welfare. Against this, in transportation regulation especially, the more widely-held a convention or rule is (such as using red lights to indicate stop) the more useful the regulation will be. We explore this tension and discuss what position regulators might take regarding automated driving.

Data Collection

It will be critical for safety, insurance subrogation, software development, driver training and numerous other reasons to record exactly where, when, and why an ADS is engaged and disengaged. Similarly, it will be important for an ADS to record where, when, and why it broke a traffic rule. Against this, such records could be invasive of privacy. We discuss these competing requirements of data collection versus privacy as it applies to automated driving.

Loading, Unloading, Curb Space, and Parking

The current levels and nature of disorder and competition for parking and stopping at any curbside will not be workable when automated vehicles need to load or unload passengers or goods. We provide a first consideration of this problem and mention related ISO standards work.

Public Mobile Robots

Automation ranging from teleoperation through ADAS to ADS has reached commercial application for last-mile delivery, surveillance, and maintenance robots. Together this body of devices, called Public Mobile Robots (PMRs), intersect motor vehicles and pedestrians at crosswalks and integrate with loading and unloading goods at the curbside. We provide an in-depth first consideration of how this might impact accessibility, traffic, safety, planning, and many other matters, and review the related ISO standards work.

There are other policy topic areas that we might have chosen to make a focus of this study. We might, for example, have investigated the role of intelligent transportation systems and their integration with automated driving; or on carriage of passengers, such as by automated public-transit vehicles or "robotaxis", and how regulation of the vehicle-for-hire sector might change absent a driver to operate these vehicles. Other relevant policy regimes that may be affected by automated driving include criminal laws regarding impaired driving, work-zone policies, and many more. We ultimately decided to pursue those listed above as the matters that are not only most pressing to the regulator, but also those most amenable to forward-thinking regulation. We recognize that government decision-makers have limited time and resources, and thus that policy change requires efficiency: pursuit of the most good at the least opportunity cost. We hope that once the changes we recommend here have been weighed and implemented (or not), there will be occasion to consider the matters we have left aside in this work.

1.2 Methodology

To prepare this study, we undertook a variety of steps.

We conducted a literature review for each of the study subtopics, drawing on material that seemed to be the most comprehensive, thoughtful, or helpful from around the world. One of the guiding principles in making our selections was to emphasize recency. Given the state of change in the field of automated driving, almost nothing published before 2018 was useful to us. The complete list of reviewed sources may be found in the Bibliography.

We relied heavily on the 2018 report from the U.S.-based National Cooperative Highway Research Program (NCHRP), *Implications of Connected and Automated Driving Systems*. This six-volume report is unrivalled for the depth and breadth of its coverage, and its wide intended audience, which includes policymakers as well as designers, manufacturers, infrastructure managers, and automotive engineers.

We also conducted a jurisdictional scan to ensure that we understood the best and latest thinking on policy and regulation of these matters. Our efforts focused on the sub-national and national jurisdictions that we considered to be the foremost legislators and regulators on these matters. At time of writing, these were as follows:

Table 1 - National Vs Sub-national Differentiation

Sub-nationally	Nationally
• Arizona • California • Florida • Michigan • New York	• Australia • Finland • the Netherlands • New Zealand • Singapore • the United Kingdom, including both England and Scotland

Our choice of jurisdictions relies on our sense of which polities are leading in some aspect of automated-driving policy and regulation. California is home to many start-up firms in the field, and has seen significant testing and deployment, as has Arizona. New York is a particularly sophisticated sub-national actor. Michigan hosts a significant amount of economic activity relating to automotive manufacturing. Florida has proven itself willing to explore new approaches to automated driving, particularly as regards insurance, liability, and the trucking sector.

Internationally, Australia, Finland, Netherlands, New Zealand, Singapore, and the United Kingdom all score well on assessments of readiness to adopt automated driving technology. This, and the ready availability of English-language documentation of their work, led to their selection.

Additionally, on the subject of sidewalk robots, we conducted a special jurisdictional scan, to account for the fact that sidewalk robots are sufficiently new and unregulated that there are very few jurisdictions that have made progress in drafting policy. Jurisdictions that have taken sidewalk robots seriously enough at time of writing are several U.S. states, including Michigan, Pennsylvania, and Virginia, as well as Estonia, Japan, and South Korea.

Finally, we prepared a survey that included questions regarding all the policy areas identified above and invited a variety of qualified experts and stakeholder representatives from a variety of fields to reply, as shown in Figure 1.

Figure 1 - Sector Specific Survey Recipients

Sector	# of Invited Participants
Automated-driving firm	15
Academia	14
Policy and Regulation	11
Government	11
Trade association	10
Transportation and mobility	9
Original equipment manufacturer	6
Insurance and liability	5
Civil society	5
Logistics	5
Law enforcement	5
Robotics	3
Grand Total	**100**

Survey Recipients by Sector

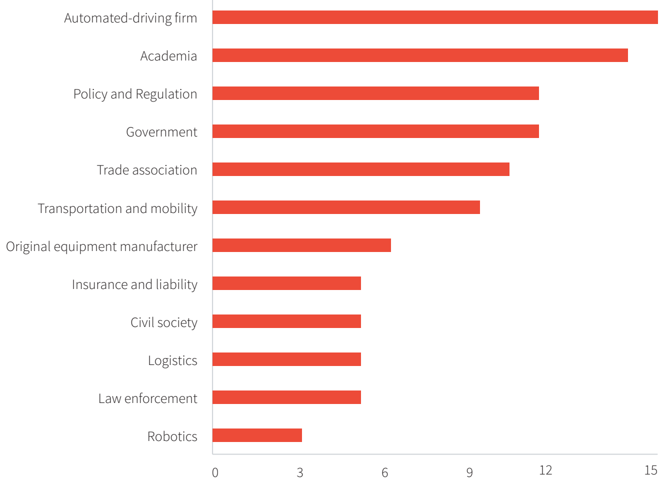

Out of 100 invited subjects, 33 completed the survey, for a completion rate of 33%.

This survey's findings are not statistically significant; the pool of responses was too small and did not represent any community. Nonetheless, we found the results revealing, and they helped to inform our recommendations. We note that unanimity was practically non-existent; for any given question, there was at best a weak majority in favour of a response, and often merely weak pluralities. From this, we conclude that the field of automated-driving regulation is still nascent, or at least most understanding of it is.

Finally, we identified fourteen experts whose body of knowledge regarding one or more of our policy topics (including sidewalk robots) was sufficiently detailed that we would benefit from a detailed discussion of the matter at hand, and conducted personal interviews with them. To ensure a full and frank airing of views, we conducted these interviews under the "Chatham House rule", i.e., without attribution. Each expert reviewed the questions asked in the broad survey, and received in advance initial questions for discussion, although these questions were meant to, and did, serve as a platform for a broader free-form discussion that allowed for delving into matters that arose over the course of conversation. The bulk of these interviews were devoted to an expert's particular area of specialty.

2

Hype, Inflated Expectations, and Disillusionment

In Gartner's Hype Cycle, the Trough of Disillusionment follows from the Peak of Inflated Expectations.[2] The Trough is where we are now regarding vehicles with automated driving systems, or 'driverless vehicles' as they are popularly called. It is a terrible time—full of criticism, despair, doubt, creaking progress, financial loss, and ridicule—exactly as Gartner's model predicts.

We have no doubt this disappointing time will be followed by the Slope of Enlightenment, wherein we will finally understand how to value and apply these innovations. Will it be about a lucky minority of people being able to safely sleep behind the wheel on their way home from work? Will it be about large numbers of people no longer needing to own a personal vehicle? Will it be about replacing lumbering municipal bus systems with smart, flexible shuttles silently and safely gliding along calm streets, driving people to train stations or shopping centers? Or will it be about everyone who cannot hold a driver's license now owning a private vehicle of their own, oblivious to compounding congestion because they will be blissfully entertained while crawling in traffic?

Whatever the answers to these questions, we believe that when the outcome is determined, the Slope of Enlightenment will be followed by a Plateau of Productivity. In that phase, whatever is finally worked out will reach scale, *regardless of the virtue of the outcome.*

When that happens or, more accurately, on the way to that happening, every governing jurisdiction that intends to permit and license vehicles that are driven by software and sensors will need to rethink how motor vehicles are governed, how their owners and users are to behave, and how all of this is to be managed and enforced. These rule-makers will need a new rulebook.

Our white paper sets out some recommendations for that rulebook.

The Gartner Hype Cycle sometimes churns slowly. Given the current slow pace of AV deployment and the regulatory, infrastructural, and user change required, we suggest that it will take another decade to crawl up the Slope of Enlightenment (2025-2035) and yet another to find our way securely to the Plateau of Productivity. If our thinking is right, then this moment of deep dissatisfaction with the capability of automated driving systems is unlikely to dissipate significantly until 2040 or later.[3]

In the interim, though, we live in a world where automated driving, however inadequate and however constrained, is active on our roads. It has already started. Unless banned, it must be regulated. We hope this white paper helps regulators in that process.

-

We have arrived at 2023, eight years after peak hype for the driverless vehicle.[4]

2 The Gartner Hype Cycle is a popular graphical representation of the lifecycle stages a technology goes through from its initial development to its commercial availability and adoption. Its five stages are: Innovation Trigger, Peak of Inflated Expectations, Trough of Disillusionment, Slope of Enlightenment, Plateau of Productivity.

3 There is no clear correlation between the SAE levels and the Gartner Hype Cycle stages. The reader should not pair the likely unattainable SAE Level 5 with Gartner's Plateau of Productivity. Gartner is merely telling us that we will eventually find the optimal and most productive way to use a particular innovation.

4 https://www.theguardian.com/technology/2021/jan/03/peak-hype-driverless-car-revolution-uber-robotaxis-autonomous-vehicle

Since 2015, we have experienced a torturous crawl of slow progress in vehicle automation dominated by the spectre of endless 'edge cases.' An alarming number of firms that made "Driverless Is Coming Soon" promises have succumbed to acquisition or bankruptcy. The steady stream of "next year" and "five years away" prognostications became a running joke.

From 2005 to 2015, the progress curve of vehicle automation appeared to rise sharply. Pundits gauged this as "exponential." During this period, the marketers of 'driverless' repeatedly promised its arrival before 2020. 2030 was sometimes identified as the essential end of private ownership with a tsunami of "mobility-as-a-service" vehicles replacing private car ownership as well as public transportation. Some authors saw a coming era of trains, automated shuttles, and taxis replacing almost everything else. Several prominent innovators in the space painted a future of almost solely private cars and taxis and the end of mass transportation.

Since the mid-1950s, artificial intelligence has developed in phases, with spurts of investor-thrilling progress ending in long doldrums. What appears to be exponential for a short period—such as we thought we were seeing in the driving automation decade leading to 2015—is replaced by something known as an "AI winter" in which progress slows to a near-halt. As shown in Table 2, this change has had an effect; the optimism of the mid-2010s has faded into significant pessimism.

What happened is simple. Until now, AI has arrived in sigmoidal or S-curves, not the exponential curves that excite marketers and futurists. Progress is real, but the sense that it is permanently exponential is illusory. The front half of a sigmoidal process looks like the beginning of an exponential process. There is no other evidence to suggest exponential progress has arrived. While we may one day achieve SAE Level 5 capabilities, if we do achieve it, we will not get there in a single, exponential sweep, as has been promised. If we get there, it will be by fits and spurts, i.e., a sequence of overlapping S-curves. This is more easily imagined as exponential progress when viewed from some distance, such as the sweep of the past half-century or the next one.

The mistake most pundits, innovators, and futurists fall into is the common error of extrapolating beyond the data we have. A failure mode of human reasoning is to rely on simple linear models for our forecasts, such that we expect that exponential progress will breed more exponential progress. Humans prefer simple linear models for their predictive certainty, but these modes fail us. Any short-term experience with artificial intelligence is typically optimistic, but our 70-year experience is that AI waxes and wanes.

Although progress in automated driving systems may have slowed or become disappointing, it is impossible to accurately predict when the next surge of progress will take place, what features it will bring, or how far the next spurt in progress will take us. There are still those who believe, without evidence, that we will reach SAE Level 5 in a decade or two, and there are those who believe we will never get there, and that Level 5 is an engineering aspiration. We will not choose sides in this white paper. Rather, we write as though we will get remarkably close in the sensible future. In the meantime, we are quite confident that we will see a decades-long rise in competence and the geographical spread of the capability of SAE Level 4. That is all we need to justify the publication of this white paper.

-

Throughout most of the past two decades, there have been many frustrated calls for governments to plan or prepare infrastructure, regulations, and local economies for the massive change to be brought about by automated vehicles. How should governments respond to these calls?

Governments are always outpaced by innovators, investors, and marketers. This is not a criticism but merely an observation of how change unfolds in complex, liberal, free-market societies. No innovation is simultaneously and immediately acceptable, adoptable, affordable, applicable, and understandable by society as a whole. That is one critical reason governance will always move far more slowly than technological innovation. Hype is merely an early, and sometimes inaccurate, warning system.

The widespread displacement of the working horse by the internal combustion engine took four decades. And it took four more decades before the associated infrastructure and private vehicle ownership saturated the developed world. The rest of the world is still catching up.

Automated driving was never going to become next year's all-at-once event. Automated driving is a gradual, generational change and will likely take more than two generations to embed and saturate. The social response to driving automation is nascent, but it has begun.

By now, it is impossible for transportation authorities to turn away from vehicle automation and its social, economic and regulatory impacts, however constrained they may be. Too much of the technology—its sensors and its collaborative, near-automated driving decisions—has escaped the lab. Partial vehicle automation already changes driving risks, challenges driver attention, impacts vehicle repair costs, reshapes goods delivery, and influences vehicle purchasing decisions and vehicle markets.

It no longer matters for government action whether pervasive vehicle automation will take three more years or three more decades to mature; it has started. Vehicle automation is simply moving slower than innovators promised. It will take longer to become pervasive and to address the safety problem, the congestion problem, and the urban livability problem. But this will not make governance easier nor excuse its delay.

We need governance and regulatory attention now to the liability question, to the digitalization of the rules of the road, to upgrades in enforcement, to updates in driver education, to the intermediate 'mixed traffic' problem, and to numerous other matters which are already lagging change.

The remaining chapters in this white paper aim to offer a constructive summary of 12 key matters related to driving automation that require governance and regulatory attention starting now and continuing over the next decade. We think this is important enough that ministries and departments of transportation that have not yet done so should create a dedicated office to prepare the required regulatory and oversight changes.

We caution that the extreme degree of harmonization appropriate among adjacent jurisdictions may make this task more difficult than expected. This is why we emphasize our recommendations for an early start, dedicated offices, and a collaborative mandate.

The long-term promise of driving automation is significant, just as its long-term threats are concerning. The hurdles to be crossed during the next three or four decades of mixed automated and non-automated vehicles will be daunting, but a wait-and-see strategy will cost us dearly.

We named this white paper *The Driverless Endgame* for three reasons:

First, it is clear that the societal outcome of the program to automate driving will be protracted, gradual, and troublingly mixed among levels and varieties of machine capability confounded by variable human behaviour. The next decades will be both complex and complicated—a long endgame.

Second, the impact of the coming multi-decade changes will alter our cities—for better or worse—in ways possibly more substantial than the changes wrought by the initial century of automobility—an important endgame.

Third, and most importantly, the societal outcome by the end of this long period will be decided by governance, whether carefully considered or left laissez-faire—a complicated endgame.

There is a fourth reason, but one we are not prepared to make sense of. These next several decades, which will decide the contribution of vehicle automation, will coincide with the decades which decide our response to climate change. It may be that these two processes will have little to do with each other. It may be that one will make the other very much more difficult. This may be the real endgame, a dangerous endgame, but we know far too little to project a direction for that interaction.

We hope you will read the following chapters in a collaborative spirit rather than as an instruction.

Table 2 - From Optimism to Pessimism

Pre-Peak Hype	Post-Peak Hype
2005–2018	**2018 onwards**
Autonomy: The Quest to Build the Driverless Car - And How It will Reshape our World, Lawrence D Burns & Christopher Shulgan (2018)	'It's a scam': Even after $100 billion, self-driving cars are going nowhere", Bloomberg (2022)
Extolled all the potential; assumed that AI was on an exponential curve rather than a set of S-curves due to edge-cases and the complexity of ODDs	Excoriates all the overpromising, over-investment, and underperformance associated with driverless technology
Did not see the full picture	Does not see the full picture
We just need to wait a short while longer (a couple of model years), allow the technology to mature, and enjoy sharing the tsunami of inexpensive, highly available, mobility-on-demand, small vehicles	The media and industry are now in a state of despair. Transition has taken too long, and the industry apparently faces a never-ending barrage of edge cases
Affordable. Clean. Easy. Equitable. Faster. Profitable. Quick. Right-sized. Safe. Uncongested	Failed promise. Lost effort. Skepticism. Accusation
Private vehicle ownership, public transportation, and parking lots would disappear everywhere You would not own a vehicle, but your trip (ride) could be private or semi-private—may be shared with one or two strangers going your way The culture of sharing was on the rise. The complexity of shopping, child seats, and camping trips was under-considered	This pessimism indicates the withdrawal of positive anticipation for any technology past peak hype Automated driving hype has turned toward Gartner's "Trough of Disillusionment," which feels more like a "pit of despair" in this case
This is indicative of the rising expectations for any newly promised technology. The anticipation for automated vehicles was extraordinarily high because the automobility problem it promised to address is extraordinarily difficult and a scourge on our cities. It promised nearly all of the advantages of secure, tailored, motorized travel and almost none of its disadvantages—a miracle without downsides	What always happens with failed-promise technology is that we recover what value we can from the innovation that has been produced (and this abounds). We find applications that may not be as exciting or as all-encompassing as was promised, but they still provide important value, require regulation, have intended and unintended consequences, and lead to the next chapters of innovation and change. Automated driving has disappointed us at the same scale in which we were over-promised, but almost the entire "history" of automated driving is still in our future

3

Language for Automated Driving

3.1 Definitions of "Driver"

One of the most disruptive challenges to regulatory language will be the definition of "driver".

- Merriam-Webster defines driver as "the operator of motor vehicle"[5]
- One motor vehicle code, which we shall not identify, defines driver as "the occupant of a vehicle seated immediately behind the steering control system"
- Another motor vehicle code defines driver as a [natural] "person who drives a vehicle on a highway"
- The SAE3016 document assumes driver to be a natural person, but describes an automated driving system (ADS) as "capable of performing the entire dynamic driving task (DDT) on a sustained basis"
  - This implies that an ADS replaces a [human] driver when the ADS is engaged, but the SAE document does not refer to the ADS as a "driver"
  - Furthermore, a vehicle entirely under the control of an ADS is referred to as "driverless", further implying that an ADS is not equivalent to a [human] driver
  - SAE3016 defines "remote driving" as the real-time performance of part or all of the DDT by a remote driver, which implies that a remote driver is a natural human; such a human driver, often called a "teleoperator", uses vehicle sensors, effectors and radio controls to "drive" a vehiclefrom a distance[6]

Regulators may need a single word or phrase to describe that responsible entity that is executing the DDT, whether that is a human seated behind the steering control, a remote human, or an ADS. As of this writing, such a responsible (or liable) entity implies one of these three forms.

The matter of this definition takes on considerable importance in those circumstances when a vehicle might be drivable by any two or three of these forms of "driving" entities. It is when driving responsibility switches between one or another of these forms that comprise the difficulty.

3.2 SAE3016 and Regulators

In this section, we present a precise vocabulary for discussing vehicles that feature automated-driving capabilities. This framework is useful for several reasons. Firstly, it discusses driving automation in terms that are most useful for regulators, by providing for unambiguous assignment of liability for the DDT to the human driver or to the ADS. Secondly, it relies only on the language in SAE3016, and eschews any language identified by SAE3016 as deprecated.

Specifically, we recommend **use of the following terms**:

- Advanced Driver Assistance System (ADAS): a system that assists a human driver to perform the entire DDT but requires active human oversight
- Automated Driving System (ADS): a system, combining hardware and software, collectively capable of performing the DDT on a sustained basis, regardless of whether it is limited to a specific operational design domain (ODD)
- ADS-Human (ADS-H): a system that does permit a human fallback driver to disengage the ADS and assume driving the vehicle when the vehicle is in its ODD

5 See https://www.merriam-webster.com/dictionary/driver

6 Taxonomy and Definitions for Terms Related to Driving Automation Systems for On-Road Motor Vehicles, J3016_202104. Retrieved from sae.org/standards/content/j3016_202104.

- ADS-Vehicle (ADS-V): a system that does not permit a human fallback driver to disengage the ADS and operate the vehicle when the vehicle is in its ODD, though outside of its ODD an ADS-V-equipped vehicle may be operated in ADAS or other manual mode by a human fallback driver
- ADS-Dedicated Vehicle (ADS-DV): an ADS-equipped vehicle designed for 100% driverless operation within a defined ODD, such that outside of its ODD the vehicle cannot operate normally (notwithstanding provisions to move, teleoperate or tow such a vehicle for maintenance, charging, or return-to-depot purposes)

This language will inform all our recommendations in this document. Hereafter, we will only use other language to describe vehicles featuring automated-driving systems, such as "L4s" or "CASE", when directly quoting other sources. Those sources include responses to our own surveys and interview questions, as it was that work that informed the development of this material.

Our use of these terms may seem idiosyncratic, given the widespread prevalence of "level language", such as "a Level 4 AV", in literature, among specialists, and especially the popular press. Nonetheless we believe our stricter use of these SAE terms exclusively is most appropriate when considering automated driving from a regulatory perspective. In the remainder of this section (and supplemented in Chapter 4) we will build a case for this view.

The SAE International (formerly the Society of Automotive Engineers) has promulgated a typology that defines "levels" of driving automation. This typology ranks driving automation systems on an *engineering-performance* scale from zero—not automated, i.e., no driving automation—through five, which means fully automated, not requiring human driving oversight within the vehicle in any driving circumstance that a human driver could reasonably handle. This SAE terminology is ubiquitous among observers of the AV field.[7]

While we recommend a strict adherence to SAE 3016-202104 ("SAE3016") terminology when referring to the engineering aspects of automated driving, we recommend **against reliance on SAE3016's engineering *levels* classifications when attempting to describe or determine human agency** or the definitions, boundaries, or conditions of said agency in matters such as driver responsibility or liability. Instead, we will rely entirely on the SAE3016 definitions for ADAS and ADS and whether one or the other is engaged.

Put another way, SAE3016 "levels" are helpful for engineers but distracting for regulators. The reason is that "levels" are correlated with available ADS capability but only engagement of ADS correlates with the assignment of liability. It does not matter if the ADS is installed within a vehicle classified as Level 3, 4 or 5; it matters whether the ADS is correctly configured and engaged.

Throughout this document the SAE 'levels' terminology may be used to describe engineering performance and features of driving automation, but we have tried to avoid its use, except when directly quoting other sources. This is because, in part:

- Many sources, including authors, jurisdictions, and other published sources, do not adhere to the SAE3016 language recommendations
- Many sources use the 'levels' (and other) language incorrectly such that the source's findings may be ambiguous; SAE3016 describes this use as "deprecated"
- The 'levels' language obscures more than it reveals; it draws engineering performance distinctions that may not be useful for regulators seeking a responsibility, liability, rules and enforcement framework

To justify this decision, we rely only on the SAE3016 document, which has drawn similar conclusions. Beginning here and for much of the rest of this subsection, we have copied or paraphrased text directly from SAE3016, pp. 34-36. We include this information, in depth, as a caution against using any non-SAE3016 terminology to describe driving automation, and against relying on the SAE3016 engineering-levels terminology to describe human driver liability and responsibility when using vehicles equipped with driving automation.

SAE3016 identifies certain deprecated terms to be avoided because they are functionally imprecise, and therefore misleading, and/or because they are frequently misused by attribution to Levels 1 and 2, in which an automated driving system (ADS) does not perform the entire DDT.

According to SAE3016, these deprecated terms include (among others): *Autonomous, Self-Driving, and Robotic*:

> These vernacular terms are sometimes used—inconsistently and confusingly—to characterize driving automation systems and/or vehicles equipped with them. Because automation is the use of electronic or mechanical devices to replace human labor, based on the Oxford English Dictionary, automation (modified by "driving" to provide context) is the appropriate term for systems that perform part or all of the DDT. The use of terms other than *driving automation* can lead to confusion, misunderstanding, and diminished credibility.
>
> The term *Autonomous* has been used in the robotics and artificial intelligence research communities to describe systems that have the ability and authority to make decisions independently and self-sufficiently. Over time, this usage has been casually broadened to not only encompass decision making, but to represent the entire system functionality, thereby becoming synonymous with automated. This usage obscures the question of whether a so-called "autonomous vehicle" depends on communication and/or cooperation with outside entities for important functionality (such as data acquisition and collection). Some driving automation systems may indeed be autonomous if they perform all of their functions independently and self-sufficiently. But, if they depend on communication and/or cooperation with outside entities, they should be considered cooperative rather than autonomous.

We note that these distinctions among a*utomated, autonomous, cooperative*, and so forth have implications for liability, responsibility, and human agency.

7 At time of writing, the latest version, which is the version relied on, here, is available at https://www.sae.org/standards/content/j3016_202104

> Some vernacular usages associate *autonomous* specifically with full driving automation (Level 5), while other usages apply it to all levels of driving automation, and some state legislation has defined it to correspond approximately to any ADS at or above Level 3 (or to any vehicle equipped with such an ADS).
>
> Additionally, in jurisprudence, *autonomy* refers to the capacity for self-governance. In this sense, also, *autonomous* is a misnomer as applied to automated driving technology, because even the most advanced ADSs are not "self-governing." Rather, ADSs operate based on algorithms and otherwise obey the commands of users.

This last point is critical to our recommendations concerning responsibility and liability, since there is always a degree of human agency involved, even if only selecting the time, start, and end points of the trip. In other words, a regulator may determine that:

1. A vehicle under ADS operation of the DDT shall be responsible for obeying the rules of the road and for safe operation of the vehicle into which it is installed, even as
2. A human user who has set the ADS to its task remains liable for other vehicle user requirements.

For these reasons, neither SAE3016, nor our recommendations, use the popular term *autonomous* to describe *driving automation*. As SAE3016 argues, while it may be possible to achieve Level 4 and 5 *driving automation*, it is, properly speaking, the case that no vehicle will ever be *autonomous*, and indeed we would not want one to be.

> The meaning of *Self-Driving* can vary based on unstated assumptions about the meaning of driving and *driver*. It is variously used to refer to situations in which no driver is present, to situations in which no user is performing the DDT, and to situations in which a driving automation system is performing any part of the DDT.
>
> The term *Robotic* is sometimes used to connote Level 4 or 5 *driving automation*, such as a closed-campus ADS-DV or a "robotic taxi," but it is technically vague because any automation technology could be considered to be "robotic," and as such it conveys no useful information about the ADS or vehicle in question.
>
> In addition to these deprecated terms, SAE3016 recommends against using *Automated* or *Autonomous Vehicle*—i.e., terms that make vehicles, rather than driving, the object of automation, because doing so tends to lead to confusion between vehicles that can be operated by a (human) driver or by an ADS and ADS-DVs, which are designed to be operated exclusively by an ADS. It also fails to distinguish other forms of vehicular automation that do not involve automating part or all of the DDT.
>
> Moreover, a given vehicle may be equipped with a driving automation system that is capable of delivering multiple driving automation features that operate at different levels; thus, the level of driving automation exhibited in any given instance is determined by the feature(s) engaged.
>
> As such, the recommended usage for describing a vehicle with driving automation capability is "Level [1 or 2] driving automation system-equipped vehicle" or "Level [3, 4, or 5] ADS-equipped vehicle." The recommended usage for describing a vehicle with an engaged system (versus one that is merely available) is "Level [1 or 2] driving automation system-engaged vehicle" or "Level [3, 4, or 5] ADS-operated vehicle."

Paragraph 7.3 of SAE3016 also cautions about the use of the word *control*:

> In colloquial discourse, the term "control" is sometimes used to describe the respective roles of a (human) driver or a driving automation system (e.g., "the driver has control"). The authors of [SAE3016] strongly discourage, and have therefore deliberately avoided, this potentially problematic colloquial usage. Because the term "control" has numerous technical, legal, and popular meanings, using it without careful qualification can confuse rather than clarify. In law, for example, "control," "actual physical control," and "ability to control" can have distinct meanings that bear little relation to engineering control loops. Similarly, the statement that the (human) driver "does not have control" may unintentionally and erroneously suggest the loss of all human authority.
>
> If "control" is to be used in a particular driving automation context, it should be carefully qualified. To this end, the one using the term "should first describe the control system they actually intend: the goals, inputs, processes, and outputs to the extent they are determined by a human designer and the authority of the human or computer agents to the extent they are not." Refer to Smith, B.W., "Engineers and Lawyers Should Speak the Same Robot Language," in Robot Law (2015), available at newlypossible.org.

The SAE3016 defense of careful terminology has been ignored or overlooked by most writers, reporters, and even many professional or academic researchers. It is because of their critical message for the regulator considering responsibility and liability, especially product liability and driver/operator liability, that we have taken the extraordinary step of including verbatim quotation from SAE3016 pp. 34-36. to the extent that we have.

By way of postscript: in a recent (2022) publication, the Scottish Law Commission elected to contradict the SAE and use the term "self-driving". The Commission wrote:

> The term "self-driving" is not used by the SAE, who describe it as a 'deprecated term'. We use the term because it can be given its own specific definition and does not carry other meanings in the SAE Taxonomy. As we explain in Chapter 3 [of the referenced document], we use it to indicate a legal threshold. Once a vehicle has been authorised [sic] as having a "self-driving" ADS feature, and the feature is engaged, the human in the driving seat is no longer responsible for the dynamic driving task.[8]

8 Scottish Law Commission No. 258 (2022) Automated Vehicles: joint report, paragraph 2.22, https://www.lawcom.gov.uk/project/automated-vehicles/

The Scottish Law Commission report chooses to ignore SAE's language advice while still relying on its engineering advice. We respectfully disagree with this choice. We believe that it is always the most reliable path to follow the standard. Even if that were not the case, we note the significant confusion that these engineering terms have already generated in public discourse. No regulator should import that confusion into its regulatory domain. We are firm in our conviction that the term *self-driving* has been rendered ambiguous for regulatory use. We recommend strict adherence to SAE terminology.

4

Liability, Responsibility, and Insurance

Deploying vehicles with automated-driving capability on public roads raises issues regarding liability, responsibility, and insurance. Failure to adequately differentiate, communicate, and regulate these matters will significantly hamper the deployment and benefits of ADS. Issues of particular concern include how liability will be assigned, how responsibility will be determined, what insurance instruments will be available, and how insurers and regulators will require and use related reporting data.

Regulators must assign clear roles and responsibilities to all the actors, provide unambiguous direction (through legislation) that delineates how liability will be enforced, and promulgate strong data-usage policies. They must further ensure that the public has awareness of all these changes.

In our view, this policy area is the one where the gap is largest between what decisions will be required and what regulators have so far accomplished.

4.1 A Framework for Determining Liability Regarding Automated Driving

This section provides a three-step framework to distinguish liability between human drivers and driving automation systems:

1. List the cases for which liability must be determined
2. Establish a precise liability demarcation between the human driver and an ADS
3. Establish a way to handle the switching lag between the liability state changes

4.1.1 The Cases for Which Liability Must Be Determined

Before we describe a framework to distinguish liability between human drivers and driving automation systems, we enumerate the limited variety of equipment circumstances, cases, and responses, involving liability using only the SAE3016 definitions for ADAS, ADS, ADS-DV, DDT, and ODD.

Some vehicles may be equipped:

1. With ADAS (only) and engageable by a human driver who is always responsible for the DDT
2. With ADS such that when engaged, the ADS is responsible for the DDT
3. With ADS such that when not engaged, a human driver is responsible for the DDT
4. With both ADAS and ADS capabilities
5. Such that the human driver *is* permitted to decide that the ADS shall be engaged
6. Such that the human driver *is not* permitted to decide where and whether the ADS shall be engaged
7. Such that the ADS will decide whether and when the human driver must take over the DDT
8. Such that the ADS never allows a human driver to carry out the DDT

In some cases, the human driver will be present in the vehicle, but in others, the human will be a teleoperator. An ADS might seek intervention from either; similarly, it might be the case that an ADS will reject requests from a human to engage. The following cases are possible but are system failures:

1. An ADS-equipped vehicle operates outside of its ODD due to an ADS error
2. An ADS-equipped vehicle operates outside of its ODD due to wrongful tampering
3. The fallback driver (see Glossary) in an ADS vehicle is non-responsive in the case that a fallback driver is part of the driving system
4. An ADS-equipped vehicle is disabled, apprehended, or abandoned, and there is neither a responsible driver nor an attentive teleoperator available

Regulators must have a set of principles on which to rely that will address all of these cases. In our view, the following set does so as efficiently as possible:

There must be no ambiguity regarding which actor is the responsible actor. In any motor vehicle, either its human driver or its driving automation system is responsible and liable in any particular driving circumstance. It must never be the case that a human driver may assume, incorrectly, that the driving automation system is liable for the DDT when it is not. Specifically, a human driver must never:

- Believe that an ADAS relieves them of ultimate responsibility and liability for the DDT
- Believe that, or behave as though, an ADS is engaged when it is not
- Engage an ADS when that ADS is not within its ODD

Similarly, it must never be the case for the provider of a driving automation system to depend, incorrectly, on its human driver to be responsible for the DDT, if they are not disposed to do so.

We explicate this as follows. To remove ambiguity regarding whether an ADS is the responsible actor, it must be the case that whenever, wherever, and however an ADS is engaged, that ADS is responsible and liable for the DDT:

1. *Either* it must be physically, mechatronically, and logically impossible for an ADS to be engaged or remain engaged outside its ODD
    - This means that an ADS, as part of its design, must protect itself from becoming engaged or remaining engaged outside of its ODD
2. *Or* the act of causing an ADS to engage or remain engaged outside it its ODD must be regulated as impermissible
    - This means either, or both, the ADS provider is liable for permitting this circumstance, or a human actor (driver or other) is liable for causing this circumstance

This last point further implies that the ADS provider must provide its own mechanism to determine and record a violation of its design parameters.

There must be no undefined or shared liability period, including during the process of switching into or out of ADS engagement. It is unworkable for there to be any ambiguity regarding liability or responsibility during engagement or disengagement of an ADS.[9]

Finally, *the teleoperator for an ADS system is part of the ADS system*. Whenever an ADS requires intervention from its teleoperator, that teleoperator accepts or shares liability with the ADS for the DDT.

In the case of a disabled, apprehended, or abandoned vehicle, the teleoperator or fleet operator accepts or shares liability for the vehicle user requirements.

Based on these principles, the following are our recommendations for implementation.

9 An article in *Government Technology* examined the question of assigning, and possibly sharing, liability between the manufacturer and driver. "[Coverage] might vary by circumstance. For cars with advanced driver-assistance features, the driver is still 'responsible for the overall operation,' so liability would fall under personal auto. If the vehicle malfunctioned, the manufacturer's product liability could be invoked — as already happens with standard cars. But even with a fully autonomous vehicle, individuals could have liability exposure if they didn't do maintenance such as installing software updates… Thomas B. Considine, CEO of the National Council of Insurance Legislators, said there aren't yet special legal or regulatory requirements related to personal insurance for autonomous vehicles. The Uniform Law Commission — a national organization of legislators and their staffs, judges, lawyers and professors that drafts model laws when states are seeking uniformity — considered creating one. But it ultimately decided insurance issues were complex and outside its scope, said Bryant Walker Smith, who worked on the idea." [Sagalow 2021]

4.1.2 Demarcation Between Human Driver and ADS Liability

The SAE3016 distinction between *Driver fallback* and *System fallback* provides a clear, unambiguous standard for driving automation. This can be used beyond its original intention to regulate for liability and responsibility. We recommend that the regulator **establish a regulatory classification standard for driving automation that recognizes Driver fallback and System fallback, exactly as described by SAE3016**.

To further illustrate the concern for strict adherence to SAE language clarity, we have taken our Figure 2 directly from Figure 12 from p. 34 of SAE3016, rather than any re-interpretation by other authors (though we have added a red box to emphasize the relevant distinction). This is directly from the standard. It is unambiguous and unencumbered by unintended nuance.

Regarding responsibility and liability, the regulator must focus on the sharp line between Driver fallback and System fallback—the second grey arrow from the top of Figure 2. This line is clear in the engineering standard for driving automation.

We offer, below, a precise automation-status delineation of (human) Driver fallback and (ADS) System fallback responsibility. This delineation refers to various SAE driving automation levels (engineering capabilities) but is not aligned with specific "levels." Instead, it distinguishes constrained user choices from a defined set of driving automation availabilities, as previously described in Chapter 3:

1. ADAS Only—No ADS
2. ADS-H: ADS engaged by a human driver, specifically the in-vehicle fallback driver, on stipulation that the ADS must prevent engagement outside of its ODD[10]
3. ADS-V: ADS engaged by the vehicle, specifically by the ADS itself or its teleoperator; in vehicles such that human drivers cannot override or prevent this
4. ADS-DV: ADS is always engaged and provides no opportunity for user operation

Note: According to SAE3016, ADS-DV is not equivalent to SAE Level 5. Rather it can encompass Level 5 and some Level 4 instances, provided that vehicles so enabled provide passengers use of the vehicle if, and only if, the ADS is engaged. As a consequence, there is no ambiguity regarding the party responsible for the vehicle's operation.

10 The reason that an ADS may not be engaged outside of its ODD, even by the owner of a personal ADS-equipped vehicle, is that outside of its ODD, there can be no appropriate mechanism to ensure that a vehicle so equipped is capable of safely managing the DDT. In addition to potential harm to the vehicle's owner, who may be the human driver or human passenger within the vehicle, there is potential harm to any other proximate vehicles or passengers in those other vehicles. In other words, there can be no property rights of the owner of such a vehicle that would override this safety concern. Hence, we recommend that the manufacturer, supplier, or maintainer of the ADS be held liable for an error condition that allowed the ADS to become or remain engaged outside of its ODD. While we admit an exception if the ADS has been tampered with, even then, the manufacturer, supplier, or maintainer of the ADS must be considered as potentially liable for not having a safeguard against such tampering. Subrogation, in such circumstances, may be complex and should be considered by the regulator well in advance of being required.

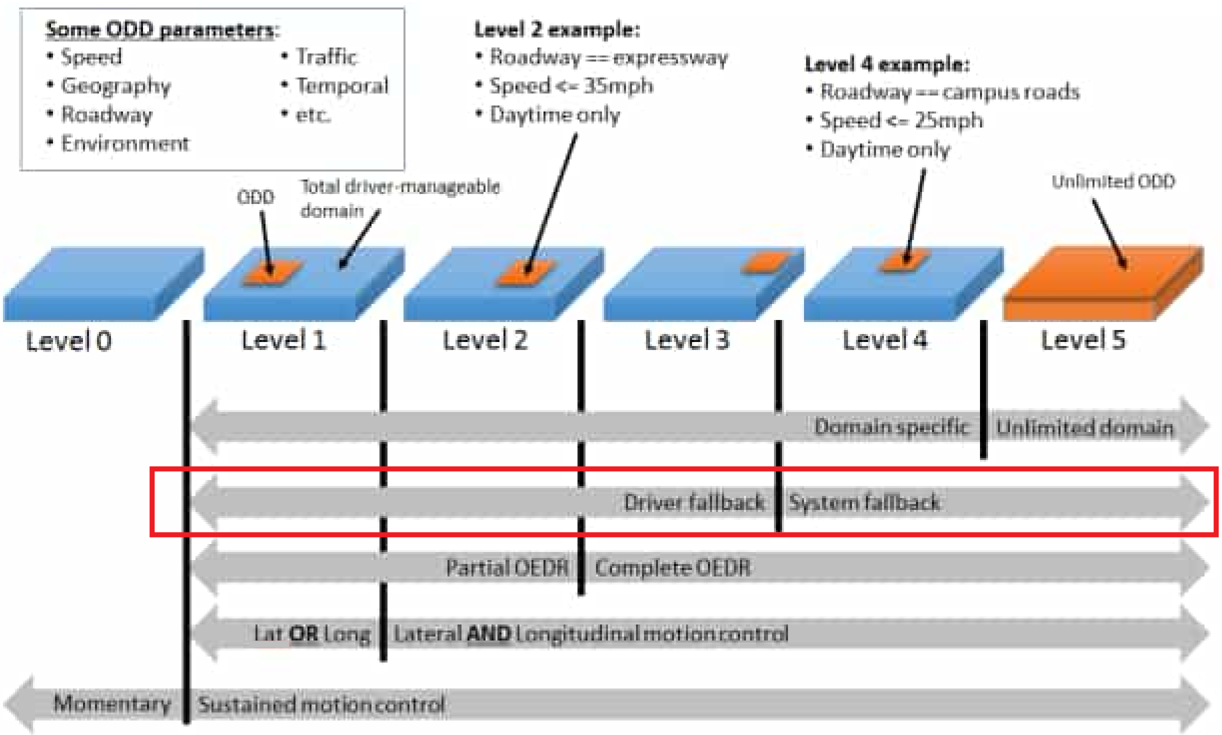

Figure 2 - ODD Relative to Driving Automation Levels, as taken from SAE3016 Figure 12

4.1.3 ADS-DV, ADS-H, and ADS-V

As a reminder, SAE3016, paragraph 3.2 defines an ADS as "… the hardware and software that are collectively capable of performing the entire DDT on a sustained basis, regardless of whether it is limited to a specific operational design domain (ODD); this term is used specifically to describe a Level 3, 4, or 5 driving automation system".

SAE3016, paragraph 3.32.3 defines an ADS-Dedicated Vehicle (ADS-DV) as "An ADS-equipped vehicle designed for driverless operation under routine/normal operating conditions during all trips within its given ODD (if any)."

From a liability perspective, ADS-DV is straightforward. An ADS-DV-enabled vehicle presumably imputes all liability (exclusive of external interventions such as force majeure, infrastructure failure, another vehicle, or criminal intervention) to the makers, maintainers, providers, and teleoperators of the vehicle enabled by such a system.

ADS-H and ADS-V present more complicated cases because a human fallback driver is present and, in the case of ADS-H, may decide whether the ADS is engaged (that choice is not available for ADS-V). Consider the baseline assumption that when any ADS is engaged, the responsible and liable party for both violations and incidents, excluding external faults, would be the provider of the system. This assumption implies that the ADS must be maintained, unmodified in a way that adheres to the ADS design determined by the ADS designer/manufacturer and certified by an applicable certification process.

This assumption also implies that an ADS vehicle of any type must be operating within its ADS ODD; hence no human driver should be able or permitted to engage an ADS outside of its ODD. This is a given by definition in the case of ADS-DV, which does not have an opportunity for a human fallback driver in the driver's seat and is expected to rely on a teleoperator to manage emergencies.

In the case of an ADS deployed in a vehicle for which a human fallback driver is present in the vehicle, we must consider the prevention of the fallback driver from engaging the ADS outside of its ODD. The reason for this critical constraint is that a human driver cannot be expected to be aware of every time, place, and circumstance defining an ODD.

For these reasons, we recommend that the regulator **prohibit any ADS that can be enabled or that can remain enabled outside of its ODD**. Here, we differ from SAE3016, paragraph 3.12, note 6: "While performing DDT fallback, an ADS may operate temporarily outside of its ODD." This note from SAE3016 does not define the meaning of 'temporarily'. A regulator cannot tolerate such ambiguity especially in regard to liability, and therefore should either regulate that any ADS provide sufficient notice to a fallback driver prior to the end of an ODD (where "sufficient" requires a precise, measurable definition), and to ensure, via some stated method, that the fallback driver has taken over before exiting the ODD. Failing that, the ADS must declare a DDT performance-relevant system failure and initiate a recovery or emergency procedure.

SAE3016, paragraph 3.12, example 3 describes the appropriate action for an ADS-DV vehicle when it "experiences a DDT performance-relevant system failure. In response, the ADS-DV performs the DDT fallback by turning on the hazard flashers, manoeuvring the vehicle to the road shoulder and parking it, before automatically summoning emergency assistance…"

We recommend that the regulator minimize the frequency of such emergency procedures by **requiring generous and specified lead-times when the ADS anticipates the end of an ODD**. How long these lead times need to be is still the subject of experiment and debate. However, in the case of an unexpected termination of an ODD, emergency procedures must be initiated with appropriate haste (again, to be defined) due to the uncertainty in the recovery of the fallback driver's attention.

There is an additional matter of ensuring that a fallback driver is not able to intervene once a recovery or emergency procedure has begun. This would be to avoid the circumstance wherein a fallback driver, aroused too late from a distracted state, would suddenly try to operate the vehicle during an emergency maneuver. We have not seen any discussion of this matter, but it requires study and regulation. No competition between the Driver fallback and the System fallback can be tolerated.

For the matter of sharply defining the liability and responsibility of the fallback driver, we recommend that regulators **divide the definition of ADS into three categories, namely ADS-H, ADS-V, and ADS-DV**, as shown in Table B:

1. ADS-H—ADS Human, which permits a human fallback driver to disengage the ADS and assume driving the vehicle when the vehicle is in its ODD; when the human driver disengages the ADS, that driver assumes responsibility and liability for the vehicle's operation[11]
   - This vehicle may always be driven by a human driver, regardless of ODD
   - This is the currently expected meaning of the ADS in a non-ADS-DV configuration
2. ADS-V—ADS Vehicle, which does not permit a human fallback driver to disengage the ADS and operate the vehicle when the vehicle is in its ODD
   - This vehicle cannot be driven by a human driver except when it is outside of its ODD
   - This would be the configuration for ADS within an ODD wherein human drivers are not permitted to operate but are still able to use non-ADS-DV equipped vehicles (such as a non-commercial vehicle)
3. ADS-DV—ADS Dedicated Vehicle, which does not permit a human fallback driver to disengage the ADS

11 In late March 2022, Mercedes announced that its ADS system, Drive Pilot, will be liable when engaged. Mercedes claims Drive Pilot operates in a limited ODD: under 65 kph, divided highways, daylight, good weather, and no sirens. This is a "traffic jam assist" system. When this ADS is enabled, the driver becomes a fallback driver, who may read, daydream, or use a phone, but may not sleep or leave the driver's seat. If the ADS needs to disengage (turn over operation to the fallback driver), a 10-second warning is provided. Notably, Mercedes claims that they (Mercedes) will be fully liable when the system is engaged, but not liable if the fallback driver fails to take over in the allotted time period. Given our definition, this is an ADS-H system, following the rule that it only operates within an ODD, it forces disengagement when the ODD comes to an end (e.g., the speed increases beyond 65 kph, it starts raining, the sun sets, or an emergency vehicle siren is audible), and the human driver can decide whether to turn it on or off, given that it is within its ODD.

Table 3 - Four Types of Autonomous Enablement, Related to Liability

	ADAS-Only	ADS-H	ADS-V	ADS-V
	ADAS decided by driver	ADS decided by human driver	ADS decided by vehicle	ADS dedicated vehicle (ADS always on)
Automation	Vehicles have ADAS capabilities that a human driver can engage and disengage.	Vehicles have ADS capabilities enabled only by the ADS or a remote operator.	Vehicles have ADS capabilities enable only by the ADS or a remote operator.	Vehicles have ADS capabilities enable only by the ADS or a remote operator.
Driver	A human driver	A human driver. The human driver is a fallback driver that drives when the ADS is disengaged. The ADS is driver when the ADS is engaged.	The ADS is the driver when the ADS chooses. If the ADS must disengage, there must be a fallback driver. The fallback driver can be a teleoperator or a human in the vehicle.	The ADS is the driver and decides correct engagement. In the case of ADS disengagement, the fleet manager or teleoperator is the driver. No provision for an in-vehicle fallback driver.
SAE levels	The ADAS could have SAE Levels 0, 1, or 2 driving automation	The ADAS combination of SAE Level 3, 4 or 5 driving automation. The ADS can be engaged or disengaged by the driver.	The ADS is SAE Level 4 driving atuomation. The vehicle has no human driver ADS switchability. In-vehicle driver controls are for special access to get to and from an ODD.	The ADS can have any combination of Level 4 or 5 driving automation. The vehicle has no ADS switchability; in-vehicle driver controls (if any) are locked away or imported for fleet operations, maintenance or emergency recovery.
Responsibility, liability	The driver is responsible, liable (except for product or infrastructure, etc. failure)	There is a finite time period between the ADS being in full control, and the human being in full control. There will be unknown (and some unknowable) interactions between the human operator and the ADS during this switch over. This switching period is the hardest to resolve...	Same as the cell to the right... ...except during special access times to get to/from and ODD, then same as the cell to the left.	ADD is responsible. Liability is a metter of subrogation among the provider cluster: • designers of the ADS (including all is sensors and effectors), • authors of its software, • mappers and maintainers of its ODD, • vehicle suppliers & maintainers

We draw this distinction between ADS-H and ADS-V in anticipation of reserving some elements of physical road infrastructure, such as specific lanes on divided highways, for ADS engagement only. Such infrastructure might be prepared or reserved on the expectation that only an ADS would operate within it. Any ADS-V- and ADS-DV-enabled vehicle could use such dedicated infrastructure so long as the human driver cannot override the ADS' decision in such cases.

We can foresee that a change in conditions (e.g., weather) might change a particular infrastructure entity from being in an ODD to not being in an ODD. If this were to happen, some of those vehicles already under ADS operation within such an ODD must either switch to fallback driver operation, exit the infrastructure, or declare a DDT performance-relevant system failure. This would temporarily render that geographic ODD from an ADS-only roadway into a mixed-mode roadway. We have not identified a solution to this, and we suggest this will remain a traffic concern for many decades.[12] In the interim, this requires further consideration.

We can also foresee the need for roadways within ADS ODDs to feature sufficient road shoulder areas to accommodate ADS vehicles that must perform the SAE3016 "DDT performance-

12 There is an obvious solution to the problem of variable ODD criteria, namely requiring common, static ODD definitions, but that solution seems to us to be technically impossible. ODDs are variable across time, place, software, incidents, weather, etc. To achieve a static, fully-agreed, all-party, ODD map is impossible, except perhaps if every vehicle was identical or if every vehicle had reached SAE Level 5—something that many now believe is impractical if not impossible. Consider, as an example, that OEM-A has an ODD limited to snow-depth SD-A and OEM-B has an ODD limited to snow-depth SD-B. A regulator would have to insist that all ADS must be constrained to SD-C, the minimum of SD-A and SD-B; and that all ADS have access to a real-time weather system that tells every vehicle on the road the exact moment at which snow depth reaches SD-C. Given what we know of both innovation and weather, it is difficult to imagine reaching such levels of achievement. For more, please see Chapter 10.2.5.

relevant system failure" maneuver (to say nothing of adequate wireless connectivity). Given that roadways in many jurisdictions may not have provisioned sufficient shoulder areas, we recommend that, in advance of licensing any ADS-enabled vehicles for regular, private or commercial operation, a jurisdiction **audit its inventory of roads to establish which ones are appropriate for ADS operation**. This implies suitable availability of road shoulders, telecommunications and proximate emergency services.

4.1.4 Switching between Liability State Changes

To reiterate:

1. ADAS vehicles always require an attentive human driver (*cannot rely on System fallback*)
2. ADS-H vehicles always require a human driver, who may engage the ADS in certain conditions (both *Driver* and *System fallback*); the human driver can disengage the ADS at will, but cannot engage it outside of its ODD
3. ADS-V vehicles typically have ADS engaged and only require a human driver in the circumstance of being out of ODD (*Driver fallback*); only the ADS can decide whether the ADS is engaged.
4. ADS-DV vehicles always have ADS engaged and make no provision for a human driver (*no Driver fallback*)

The first, third and fourth cases are relatively straightforward for a regulator to oversee, as the question of where liability lies is evident, namely with the human or the ADS, respectively. It is the second case that poses a problem: any ADS designed to be switched from *System fallback* into *Driver fallback* for any reason, including when that 'driver' is a teleoperator, blurs the unambiguous assignment of responsibility or liability.

In Table 3, the ADS-H cell on the bottom row represents the key problem that occurs when the switchover is from ADS to fallback driver, i.e., disengagement of the ADS, especially under urgent conditions.[13]

In these cases, the regulator must require a defined and unambiguous indication of the vehicle operator at the moment of a crash or traffic violation or the moments leading up to such events. Such definition and unambiguity will be extraordinarily difficult to achieve, given the time variability and unreliability in re-capturing human attention in a time period implied by "the moments leading up to a crash." For any system that is switchable between Driver and System fallback, there will be a multi-second—and possibly tens of seconds—time gap during which reassignment of liability (to a human) would be fraught.

The problems that arise during this gap include:

1. The ADS being able to secure the attention of the fallback driver
2. The quality of that human attention
3. The capability of the human to take over in circumstances sufficiently complex such that the ADS requires intervention

Human attention-switching and control recovery are well-known as insidious problems that have not yet been satisfactorily solved for human-vehicle cooperative control systems. The two best solutions are full human driving attention and full driving automation. While both are imperfect, they offer the advantage that it is clear how to assign responsibility. For the intermediate systems that permit or assign a switch from human operation to system operation and later demand a switch back from system to human, we need to address the assignment of responsibility during the time in which the switch is occurring.

For this, we recommend **a clear event-line of responsibility for ADS-equipped vehicles to remove ambiguity**, such that:

1. During the switch from human to ADS, the human is responsible until the switch is complete as determined and signalled by the ADS (see Figure 3 following)
2. During the switch from ADS to human, the ADS is responsible until the switch is complete as determined and signalled by the ADS
3. There is a uniform way to record those state switches that is reliable (often at sub-second intervals) for crash investigation and usable in the event of hearings regarding traffic violations

The first two points above require that the human driver cannot force an ADS to operate outside of its ODD and that *the ADS decides whether it is within its ODD*. These are necessary starting points because it is untenable to rely upon humans to understand when the vehicle they are using is within its ODD.

Further, it is untenable for regulators to be unable to discern which driver (human or system) was operating the vehicle during a crash or a violation, as humans would tend to blame the ADS and manufacturers of ADS would wish to blame the driver, depriving regulators of lines of clear accountability. Multiple crash events prior to 2022 appear bear out this assertion.

Methods to determine and record dis/engagement events need to be defined, standardized, and certified.

Transport safety regulators should engage with international standards for data recorders and data storage requirements and the potential need to regulate and design tests for upgraded Event Data Recorder technology. Such upgraded standards must address the need for an unambiguous way to determine that the switch from ADS to human has been satisfactorily accommodated by the human. If the ADS cannot be certain (a term that demands definition adequate to the context) that the fallback driver has successfully begun to carry out the DDT, there must be a way for the ADS to safely

13 There is an obvious solution to the problem of variable ODD criteria, namely requiring common, static ODD definitions, but that solution seems to us to be technically impossible. ODDs are variable across time, place, software, incidents, weather, etc. To achieve a static, fully-agreed, all-party, ODD map is impossible, except perhaps if every vehicle was identical or if every vehicle had reached SAE Level 5—something that many now believe is impractical if not impossible. Consider, as an example, that OEM-A has an ODD limited to snow-depth SD-A and OEM-B has an ODD limited to snow-depth SD-B. A regulator would have to insist that all ADS must be constrained to SD-C, the minimum of SD-A and SD-B; and that all ADS have access to a real-time weather system that tells every vehicle on the road the exact moment at which snow depth reaches SD-C. Given what we know of both innovation and weather, it is difficult to imagine reaching such levels of achievement. For more, please see Chapter 10.2.5.

Figure 3 - Liability State Changes for Various Types of ADS

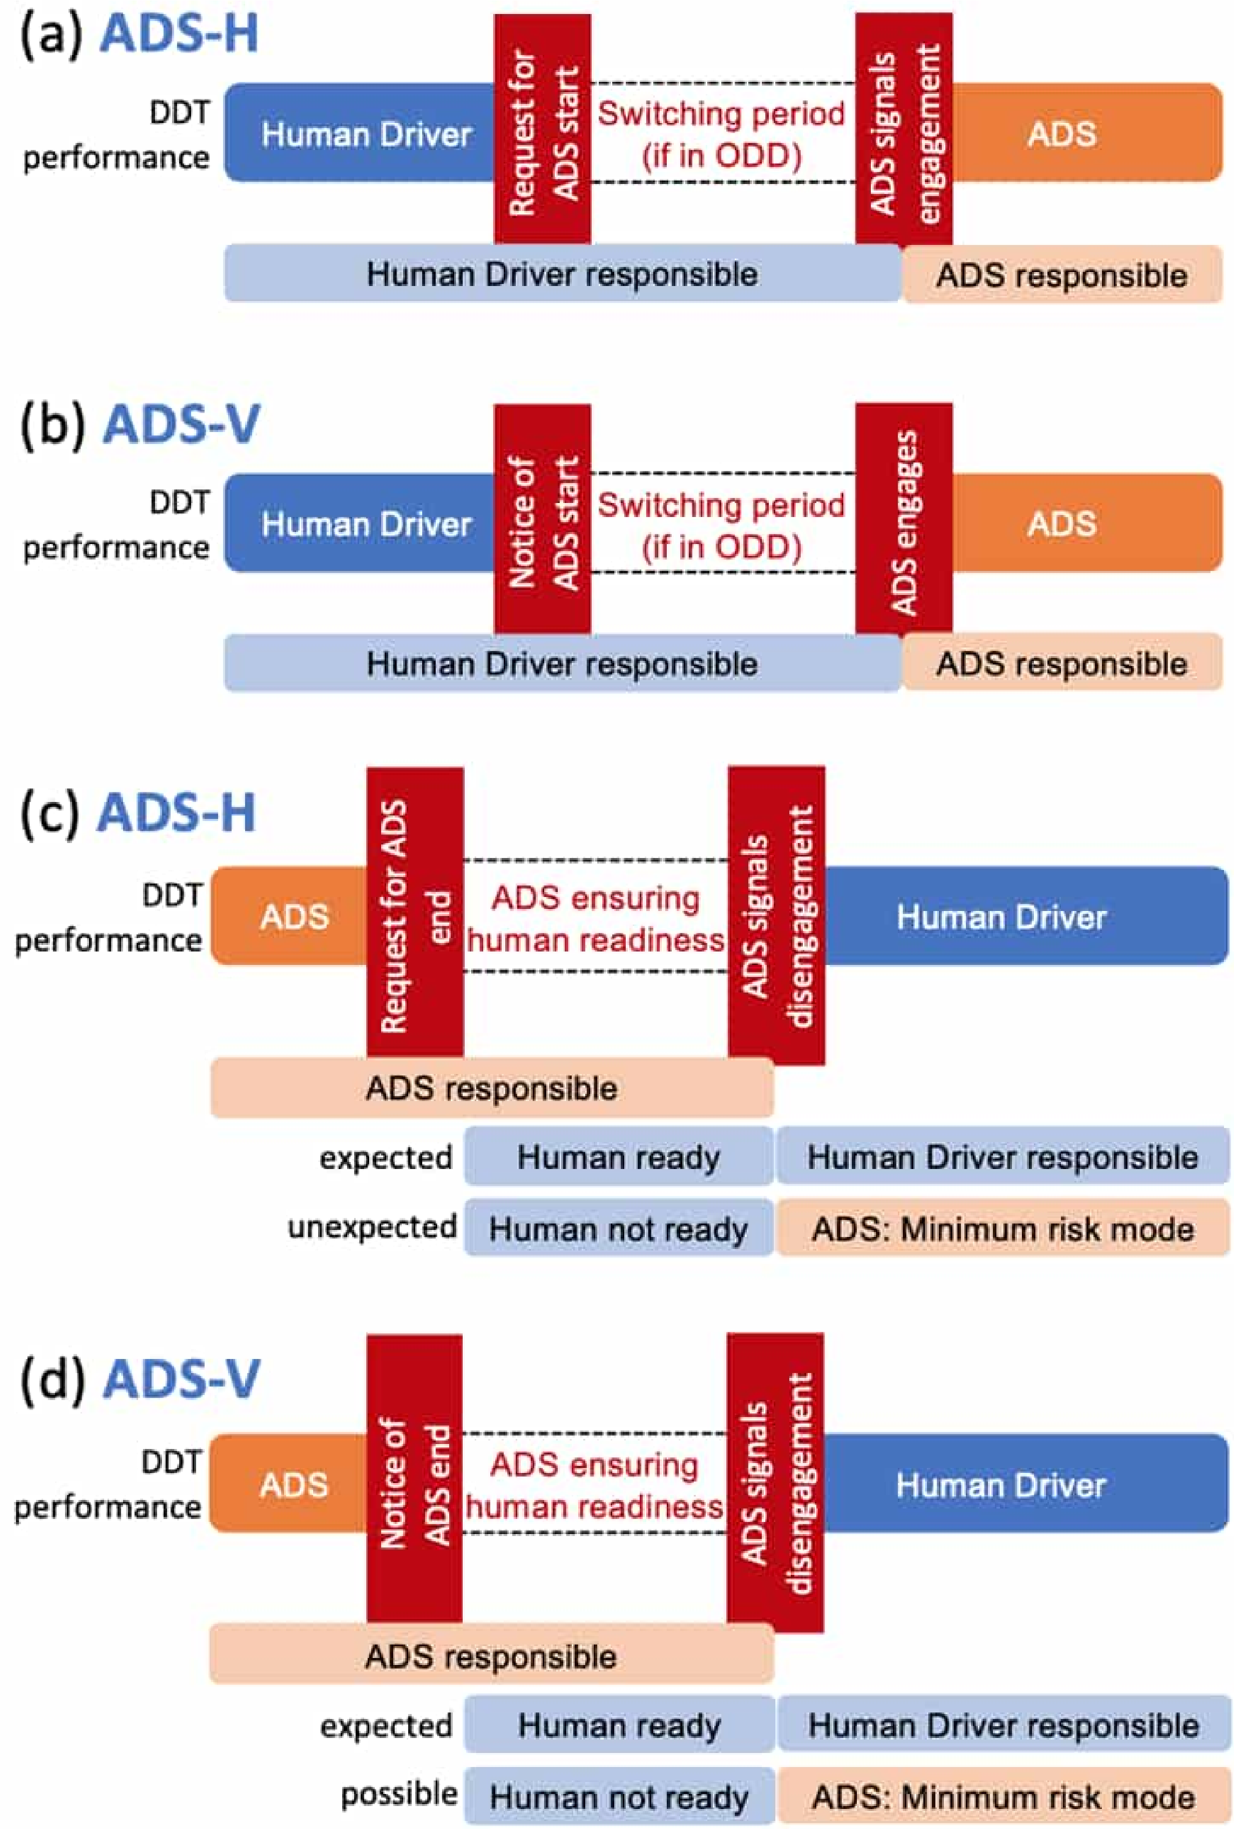

bring the vehicle to a *minimal risk position* by either bringing the vehicle to an appropriate position, presumably off the road and out of traffic or to initiate an *emergency procedure* (SAE3016 paragraph 8.6).

We add that the questions of who is to have access to this data, and when, are also fraught (see Chapter 12 for more).

In most cases, circumstances in which a switch from ADS to the fallback driver is necessary are preceded by foreknowledge of the impending end of an ODD; for example, the vehicle is about to leave the geography of the ODD, or weather-related conditions indicate the fallback driver must take over shortly. In cases of such high predictability, the problem is to have the ADS ascertain driver readiness in order to either transfer the DDT to the human driver or instead come to a safe stop. It would be the responsibility of the ADS to signal sufficient lead time and to understand how to determine a safe stopping method (minimum risk position). Failing in those circumstances would be the responsibility of the ADS.

There are harder cases in which the lead time available to the ADS may be very short due to an unanticipated end of an ODD—e.g., a recent and unmapped rockslide, a fallen tree, or a crash has occurred on the route—and the ADS cannot secure the driver's attention within the critical period of switchover. The ability of the ADS to come to a safe stopping place might be curtailed, and circumstances may arise that lead to a crash or violation. While the ADS may be considered responsible because it was operating the vehicle, it is possible that the circumstances would exonerate both the ADS and the fallback driver. Such circumstances and their potential outcomes need more consideration but can be expected to parallel cases of sudden infrastructure failure that may happen regardless of driving automation.

Hence, one potential is that any ambiguity about which side of the fallback line the driving automation system was on means assigning liability to the human owner/operator/ driver by default and that any product or infrastructure liability be determined during the subsequent subrogation process. This would put the onus on the insurance industry and the vehicle owners that pay premiums and submit claims rather than on the regulator once a framework and systems to measure switchover are fully determined.

A key concern when assigning responsibility or liability is to ensure that its definition would survive a test in court. One of the tests for product liability "looks to the safety expectations of consumers."[14] This is a key reason this white paper has avoided using the term "self-driving"; depending on the audience, the term may suggest ADS or ADAS, despite the significant disparity in capability between them.

In such a case, the difference between marketing language that sets customer expectations, and precise language in a seldom-read user's manual, may contribute to the deaths of users who overestimate the reliability of an ADAS. Language sets expectations that, in this case, are not only dangerous but may form the basis of a legal defence. From a regulatory and legal perspective, these ambiguities must be avoided.

14 Bryant Walker Smith, "Lawyers and Engineers Should Speak the Same Robot Language", in Robot Law, ed. Ryan Calo, A. Michael Froomkin, and Ian Kerr (Elgar Publishing, 2016).

4.2 Do Not Unnecessarily Relieve the Vehicle Consumer of Liability or Responsibility

In a pre-ADS world, the public generally blame most road incidents on human behaviour: inattention, intoxication, recklessness, or other human fallibilities. We expect that this will change and that the public will place the blame for machine error—a failed sensor, a software edge case, or automated brakes that failed to activate in time—on an amalgam of the OEM and the regulator. This may seem perverse, given that the party apparently most capable of preventing such outcomes is the OEM. But no matter the actual degree of fault or negligence on their part, the regulator will be perceived as an overseer, an enabler, as the party that permitted manufacturers to make the promises that were made and to sell the vehicles they do—and as the party responsible for the condition of the infrastructure that a vehicle's systems relied on.

It is a natural tendency for humans that can deflect responsibility to do so. The risk to the regulator is that at least some of this will occur, more so for any circumstance in which responsibility can be construed as ambiguous. Any unnecessary transfer of liability away from individuals (drivers, owners, operators) and toward the manufacturer and regulator creates an externality where individual self-concerned humans are more likely to behave carelessly.

This effect will be particularly pronounced during the coming period of mixed driving. For example: if I am driving a Level 0 vehicle and I crash with another Level 0 vehicle and the driver of the other car was at fault, then my recourse to recover damages is clear. This would involve our two insurance companies and perhaps a traffic citation for the other driver.

If, however, I am driving a Level 0 vehicle and I crash with a Level 3 to 5 vehicle with ADS engaged that was correctly operating in its ODD, the matter of settlement would become much more complex, and matters of proper testing of the ADS, its proper configuration, or bad vehicle design, or whether that vehicle design complied as it should have to regulations, or whether those regulations were indeed adequate, or whether its owner/operator had maintained the vehicle, etc. may have to be ascertained.

Therefore, we recommend that regulators **bias liability toward the owner/operator of the vehicle to the degree possible**. Regardless of the level of automation, the vehicle was purchased and is being operated for the purpose or profit of the person or entity that purchased and operates it. Liability must first rest with that person or entity and their insurance. Only in the case that an ADS is knowably the sole operator of a vehicle should liability be shifted away from the human driver.

As discussed in Section 4.1.4 and its Figure 3, the manufacturer of the vehicle that incorporates ADS should assume liability for ADS-avoidable violations, crashes, or injuries when the vehicle at fault is under ADS operation in a way that exonerates the fallback driver from incorrect use or other negligence. Even then, in order that the manufacturer be assigned total liability, all other parties to the incident (operator, owner, maintainer, infrastructure provider) must be blameless, which is likely to become an increasingly rare circumstance.

Establishing liability with the human driver, whether acting as primary or fallback driver, leaves the task of reapportioning that liability to the insurance industry during its subrogation processes. We see no reason to burden the regulator with that responsibility. The insurance industry will identify the safest vehicles as it sets premiums. The consumer of the vehicles, whether for personal or commercial use, will necessarily become aware which vehicles are safer or at least which vehicles carry lower insurance premiums. By letting insurance premiums "pick the winners," the consumer's purchasing decisions will have the maximum impact on vehicle safety. Government should regulate this process to ensure a common approach and an optimal (or fair) outcome among insurers and insured, as it currently does for conventional vehicles.

This approach will also mean that private consumers will (or should) be especially concerned about the safety records of vehicles they might select for family use, as would commercial purchasers acquiring vehicles for a service fleet. In this way, regulators would leverage purchasing decisions about personal and fleet machines to privilege ADS designs that exceed safety and reliability standards. It may be the case that by publishing *reliability information* as implied by insurance premiums, which themselves might be regulated to be appropriately aligned with insurance settlements, a regulator would accelerate this process.

4.3 ADS-Equipped Vehicles Need a New and Contractual Registration Process

As developed in Section 4.1.3 and its Table 3, a key recommendation for governing ADS is to define methods of ADS deployment such that there is no ambiguity regarding who or what is operating the vehicle—responsible, liable, and accountable.

There are sharp distinctions between ADAS-equipped and ADS-equipped vehicles, as shown in Table 4. As discussed in the previous section, liability rules for the driver/operator of merely ADAS-equipped vehicles, in which a human driver is obliged to perform the dynamic driving task (DDT) but receives assistance with elements of that task, should not change from vehicles that are not automated in any way. In all such cases, the driver remains responsible for the DDT. When considering ADS-equipped vehicles, we are describing a system that does not need driver monitoring *while it is engaged* but always requires fallback driver readiness. In the extreme case of ADS-DV, the required fallback driver is a teleoperator, but in all other cases, the fallback driver must be in the driver's seat of the vehicle in question. That fallback driver must be in a state of sufficient readiness to take over the DDT.

It is critical to observe that in the case of vehicles featuring either ADS-H or ADS-V, when the ADS is not engaged, the human driver is fully responsible for the DDT. This remains true

Table 4 - Types of ADS-Equipped Vehicles

ADS Class	Engagement	Disengagement	Driver	Use Case
ADS-H	The ADS is engaged by the human driver	The ADS disengages from incorrect use, which might include the vehicle being prompted to leave its ODD, or is activated outside its ODD	The driver becomes the fallback driver	This vehicle requires a human driver
ADS-V	The ADS is engaged by the vehicle; the ADS always determines engagement/ disengagement	The ADS will not disengage within its ODD A human driver is not permitted to operate the vehicle within its ODD A human driver within the ODD cannot override ADS decisions, except in cases of emergency In such cases, the vehicle may only safely remove itself from service	The user must be prepared to be the in-vehicle fallback driver in all cases that the ADS demands (e.g., end of the ODD)	This vehicle would be suitable as a passenger service vehicle in a fleet that needs occasional resort to human drivers to fill gaps in a service ODD This vehicle might not be suitable for private ownership due to its inflexibility and fickle demands on the driver
ADS-V	The ADS is always engaged	The ADS disengages only in cases of emergency override, which would take the vehicle safely out of service	The fallback driver is a teleoperator and is not in the vehicle	This vehicle would be suitable as a service vehicle for passengers or goods in an ODD without service gaps This vehicle might not be suitable for private ownership due to its inflexibility

irrespective of any ADAS capabilities that may still be active when an ADS is not engaged. By definition, such a driver is required to actively attend to the DDT.

We recommend that the appropriate regulator always establish the terms under which an ADS-equipped vehicle may be registered. The necessary and sufficient conditions for registering an ADS-equipped vehicle are as follows:[15]

1. The ADS handles 100% of the DDT, including following all rules of the road when engaged
2. The ADS always self-determines whether it is in an ODD before it engages—the ADS prevents a driver or fallback driver from engaging the ADS outside of its ODD
3. When the ADS engages, even if it does so in error, it always assumes full responsibility as though it were in a suitable ODD
4. If a vehicle driver/operator elects to engage the ADS outside of an ODD, it is the responsibility of the ADS to refuse such engagement
   - In our view, consistency requires that were a human user to somehow engage the ADS outside of its ODD, the ADS manufacturer should still be held responsible or liable for this engagement error. This means that human users may not be permitted to so engage. This also reinforces our recommendation that no unauthorized modification be permitted, discussed further in Chapter 9
5. When the ADS disengages, it must always secure the fallback driver's attention according to a defined, standardized procedure
6. This procedure must be explained to the user or purchaser of any vehicle that permits user engagement of the ADS
   - It is not suitable to sell, lease, rent, or lend an ADS-H or ADS-V to a driver who is unaware or incapable of the procedure to take over from an ADS
   - This also implies that an ADS-H-equipped vehicle cannot be used as a driverless service vehicle; this is one of the important distinctions between an ADS-H and ADS-V—an ADS-V could be deployed as a for-hire vehicle within an ODD, whereas an ADS-H could not be so deployed
7. The manufacturer, seller, lessor, or renter must declare whether the vehicle is equipped with ADS-H, ADS-V, or ADS-DV, and this declaration must be made clear to the purchaser or intended operator/user of the vehicle
8. The manufacturer, seller, lessor, or renter must guarantee that the ADS cannot be engaged or remain engaged outside of its ODD
9. In the event that an ADS-equipped vehicle is somehow operated outside of its ODD, the manufacturer is the liable party
   - The principle at work here is that ADS technology is too complex to rely on a human driver' ability to always correctly identify ODD conditions, and the ADS technology must be sufficiently robust to self-determine its ODD
   - This implies that the manufacturer of an ADS must ensure that no method of tampering can cause an ADS to operate outside of its ODD
10. The vehicle owner agrees that after registration, no seller, lessor, renter, or vehicle maintenance operator may modify any part of the ADS, including any of its software, sensors, effectors, connectors, or connections
11. The vehicle owner further agrees that in the event of a repair, alteration, or recall of a registered vehicle, only a maintenance provider approved by the party or parties that are liable for the ADS when the ADS is engaged can make the required adjustment(s)

Such a system implies that any form of marketing, sales, or performance promise that exceeds the capability or behaviour of a vehicle to be registered in a jurisdiction would be pernicious. For that reason, we recommend that the regulator **consider how to prevent the marketing of an ADAS or ADS vehicle with language that misleads regarding the extent of its capabilities**. As noted elsewhere in this white paper, vehicles with ADAS should not be marketed as "full self-driving."

It further requires that any prohibited tampering, misleading marketing, prohibited modification, repair, or alteration of a registered vehicle be regulated as a criminal offence, a recommendation that we will discuss further in Chapter 9.

4.4 Testing for Conformance with Registration Conditions

Typically, it is the responsibility of the senior-most level of a government to set standards for vehicle design, and the necessary and sufficient conditions for operations may need to be updated accordingly.

Because of the complex nature of ADS, the vast number of edge cases that will arise, and the nature of software updates that are likely to be frequent and asynchronous, it is unlikely to be feasible to test every model of vehicle and every update to its sensors and software in order to ensure compliance with necessary and sufficient behaviours and capabilities.

To avoid this expense, we recommend that the regulator **require OEMs to execute a contract that outlines the necessary and sufficient conditions for registration and require evidence of sufficient insurance** (related to fleet size) **to permit operation**.

We note that enforcing such contracts would likely require continuous processing of ADS dis/engagement data to ensure that ADS vehicles are operating within the agreed or enforceable parameters. Further, such data processing could permit a governing body to address other concerns, such as algorithmic bias stemming from input data. For all of its promise, however, requiring and regulating government access to this data is not a simple matter. Please see Chapter 8 for more on this subject.

15 "Necessary and sufficient" requires qualification. We have not included any requirement for data recording and collection. For more on this, see Section 12.1.

4.5 An Ancilliary Observation Regarding Liability for Infrastructure

In some jurisdictions, roadway infrastructure is currently identifiable as a liable contributor in a crash and for related injuries. This typically relates to the condition and design of the roadway, its signage, markings, guard rails, traffic signals, construction sites, etc. Current practices and precedents for assessing and subrogating such liability should not be altered for vehicles that may engage ADAS. This is consistent with our recommendations that the vehicle driver/operator be fully responsible for the conduct and operation of the vehicle regardless of the nature of ADAS enabled. The driver of an ADAS-engaged vehicle must attend to road conditions and irregularities, as before.

This practice may require re-examination as ADS emerges. Road infrastructure, which includes signs, markings, and traffic signals, interacts with the physical aspects of the road and its safety design. By way of illustration, we imagine a sign showing a speed limit set too high for the road curvature. We expect ADS to rely on telecommunicated information at least to some extent (the inputs from their sensor arrays that observe local conditions and circumstances providing most of the system's operating data). If an ADS receives inaccurate information via such a communication system and its IoT components, then the provider and maintenance operator for that communication system and its data would appear to be exposed to some liability. We described and recommended a standard for the management of such digital information, called METR, in Chapter 5.2[16] Because METR (or equivalent), when deployed, would fundamentally serve as the core construct for a future traffic management system that adopts that standard, such systems must be in the purview of the regulator. That purview would be at least regional and possibly national. Any error in such a system that could be shown as a cause of a crash could carry liability for the party maintaining a METR-compliant system.

One can argue that automated driving must be at least as competent as a human driver. Therefore, current practices and precedents would still hold, at least in terms of physical infrastructures, such as the road surface, bridges, guard rails, and construction sites. An ADS should recognize these conditions and take appropriate action. However, this simple assumption is currently untried. The interaction of ADS with road infrastructure at scale may reveal ADS qualities that are not anticipated. We believe that, at this time, this question cannot be answered with certainty.

In any case, we note in passing that the question of government liability for infrastructure may mean no change for ADAS, would likely mean some change for ADS-H and ADS-V, and would almost certainly mean critical changes for ADS-DV.

---

16 This refers to ongoing (unpublished, approved work item) ISO TC204 WG19 AWI TS 24315.

5

Digitalizing the Rules of the Road

5.1 At Issue

"Rules of the road" refer to the set of laws and regulations that govern the movement of vehicles on public roads. These rules cover a range of topics, regarding vehicle operation, such as behaviours at intersections and within road lanes, signalling intentions while entering and leaving roadway elements, and what to do in the event of an accident.

Until now, these rules-based behaviours have been mediated by licensed drivers who are responsible and, in most cases, held liable. There are at least two schools of thought regarding these rules in the context of automated driving.

Enthusiasts point to the fact that sophisticated ADS will have superhuman capabilities to react to road conditions, but these behaviours—speed limits, following distance requirements, expectations for traversing intersections, and more—that are calibrated to human abilities would not permit such capabilities to be used to their full potential.

Conservatives might remark that these rules will continue in force for many decades while human-operated vehicles remain on public roads. These rules comprise a large body of detailed guidelines and descriptions including tolerances and measurements. They also include phrases that apply specifically to human cognitive executive function, such as with *due care* or *safe distance*. It is too early to say if software engineers can translate all these terms into code that a robot can use. It is also conceivable that translation of concepts such as with *due care* or *safe distance* into rigorous constructions with application in courts or settlements may, in conjunction with real-world conditions and events, result in unintended consequences.

While such a body of code has not yet been optimized, we know it is being attempted and we believe that it will be achieved. Doing so, however, will likely iterate the problem into a new form, as robots executing this optimized code interact with human drivers that may respond with frustration, caution, confusion, or abandon.

It may be that the rules of the road should not be changed, or should not be changed yet, but they certainly must be translated. We hew closer to a conservative than an enthusiastic position, and as such will argue that in this translation, the sets of rules for humans and robots must be, if not identical, at least close enough so that the interaction between robotic and human behaviour is sufficiently matched to understand how mixed traffic will operate.

5.2 Policy Analysis and Recommendations

Our first and most critical recommendation is that **governments responsible for the rules of the road refine them such that they are identical for humans and ADS**. To date, the requirement to follow motor vehicle regulations, barring exceptions during AV trials and pilots, has been expressed in terms of driver-centric responsibilities.

Rules of the road should be expressed in vehicle-centric terms to minimize ambiguity. Rules as currently expressed assume a human driver, incorporating expressions such as "eye contact" or "due care." Rules must be rewritten to neutralize this language. Consider replacing specific terms: "steering assemblies" rather than "wheels" and "braking systems" rather than "pedals." Preparing two separate codes, one for human drivers and one for machines, would inhibit uniformity that would, in turn, breed confusion and create a potential for harm. There should be one code that applies equally to all driving agents. The expression of these rules must be agnostic to how agents execute those rules.

Secondly, we recommend that governments, when updating the rules of the road to address connected vehicles and ADS, **develop the master source for these rules as the digital version**.

We assert this by making several assumptions:

1. There will be fleets of vehicles that are equipped with driving automation capabilities
2. These fleets will grow in both absolute and relative size (vehicle counts)
3. For the foreseeable future, roadways will have vehicles equipped with a variety of levels of driving automation, with many operating in proximity to each other, i.e., without physical barriers for separation
4. In a further future, the overwhelming majority of ground-based passenger and goods vehicles will rely almost entirely on driving automation
5. The time at which there are vehicles that never require direct, hands-on human judgment and operation is in the very far future, and indeed may never arrive

These assumptions imply that in the near future—certainly for next two decades—only a minority of vehicles will operate entirely with digital driving rules embedded within their software systems, and in an environment with fully telecommunicated traffic control. Alongside this minority, most vehicles will continue to be operated by drivers, whether fully engaged or as backup drivers, and these drivers will continue to require rules and training. These rules and training will be available in analog forms, such as legislation published in human-readable form; bylaws; and training manuals and courses. Furthermore, such drivers will continue to require analog traffic-orchestration controls in the form of marked driving lanes, traffic signals, signs, speed cameras, and other forms of guidance legible to humans.

Given these assumptions, we predict that as driving automation comes to dominate, we will migrate from a world wherein predominantly analog rules and control systems will be replaced by predominantly digital rules and control systems. This replacement will be gradual; again, on the order of a few decades.

We have entered a long period in which transportation technologies are moving from analog to digital primacy. The end of that period, some decades from now, will be marked by traffic designers, legislators, and engineers rendering the human analog from the primary digital rather than representing the primary analog using digital means, as we do now.

By this, we do not mean merely that the rules of the road should be machine-readable (though they must be). Instead, we mean that in the future we will place analog lights, lines, and signs where a digital map indicates, instead of creating a digital map based on where the signs, lines and lights are. The relationship between analog and digital will be reversed. Legislation would begin life as a set of APIs for ADS use, and only then be translated—likely by machines—into human-readable prose. Early releases of AI language models such as ChatGPT that can turn human prose instructions into Python code and code into human prose indicate the shape this future is taking. [17]

Today, any digital representation of a system of traffic rules, traffic control, or driver training is created as a digital twin of its existing analog counterpart. The digital is required to conform to the analog. Any change in the current analog versions of driving rules, traffic control systems, training courses or licensing systems, implies that their digital representations, where they exist and where they need to codify what is in the analog world, must be updated. Google Maps is updated based on what's on the road, not the other way around. But as transportation systems become more complex with more forms of vehicles and more automated driving capabilities of those vehicles, the task of designing and governing this space will demand greater digitalization. This, in turn, will make the task of deriving the digital from the analog—along with its alignment and enforcement—increasingly difficult, hence the reversal.

In this regard, consider the ISO project M*anagement of Electronic Transport Regulations* (METR).[18] To quote from the project overview:

> We live in a transformative age for transportation. Increasingly, we are seeing pedestrians equipped (and often distracted by) smartphones, new modes of personal travel, such as e-scooters, an increase in the use of delivery services through the advancement of technologies that match drivers and travelers, and the introduction of automated vehicles whether they might be designed for personal travel, shared travel, or delivery of goods. Each of these changes in the transportation environment results in the need for new regulations, often all being applied within the same travel space.[19]

As the passage suggests, it is not merely automated driving systems that are applying pressure on transportation governance. Change and complexity come from multiple directions. Given the expected scope and pace of change, it is more efficient for a transportation authority to make the rules of the road digital-first. Digital representations of the transportation environment and its regulations are far more easily managed, updated, transformed, standardized, published, broadcasted, and distributed than are analog representations. It is also easier to build and maintain digital systems of representation and to then derive analog instances from them than it is to maintain-and-derive analog masters into digital representations. Of course, matching physical, analog infrastructure from digital instructions is not trivial. New systems of care will be required to position traffic signals, signs, and lines to precisely map the digital representation. So, while this will neither be a swift nor a one-way process, we insist that what is on the ground, in the rulebooks, in driver training courses, in our enforcement systems, and inside ADS code must match.

17 There would be human oversight and intervention throughout such processes… at least at first.

18 The project's home is located at https://iso-tc204.github.io/iso24315p1/index.html.

19 Vaughn, K., Marousek, J., (2021) METR Overview, https://iso-tc204.github.io/ iso24315p1/METROverview.pdf.

Thirdly and finally, we recommend that any road transportation authority **consider the intention and direction of the ISO METR project**:

> METR will facilitate an overall approach and the mechanisms for the distribution of road rules, where the consumers of these rules can be assured of the authenticity of the regulation or ordinance. This notion of authenticity is the critical point: if road rules are to be used to guide the operations of vehicles on roadways, the users of those rules must have confidence in the veracity of those rules. There must be a mechanism whereby the receiver of the rule can verify that the rule is correct and legitimate.[20]

METR as a concept applies to the entire life cycle of a traffic regulation beginning immediately after that regulation is codified, and includes the complete distribution chain of the regulation, from the point when the regulation is created, to when it is disseminated and shared with end-user devices. As such, METR may incorporate or define data models, distribution mechanisms, cryptographic processes, implementer roles and responsibilities and any other item or concept that is required to distribute machine-usable road regulations to users. METR's implementations will have to respect variances in scope, operations, end-user distribution mechanisms and other constraints among its participants.

Whether or not the METR approach becomes the universal solution, becoming familiar with its tenets and methods will assist regulators in their efforts to assert the primacy of digital governance of road rules, traffic control systems, enforcement, and other aspects of ground transportation governance.

Figure 4 is an indication of the current intended schedule for the development of the METR project. Any transportation authority intending to grapple with the digitalization of its rules should take a proactive measure of this work well before 2025.[21]

20 ISO (2021) METR Vision, https://iso-tc204.github.io/iso24315p1/METRVision.pdf.

21 This schedule shown in Figure 4 is an internal ISO/TC204/WG19 meeting document. It is not a formal ISO promise and is subject to change. It is used with permission of the workgroup members who authored that meeting document.

Figure 4 - Tentative Deployment for the METR

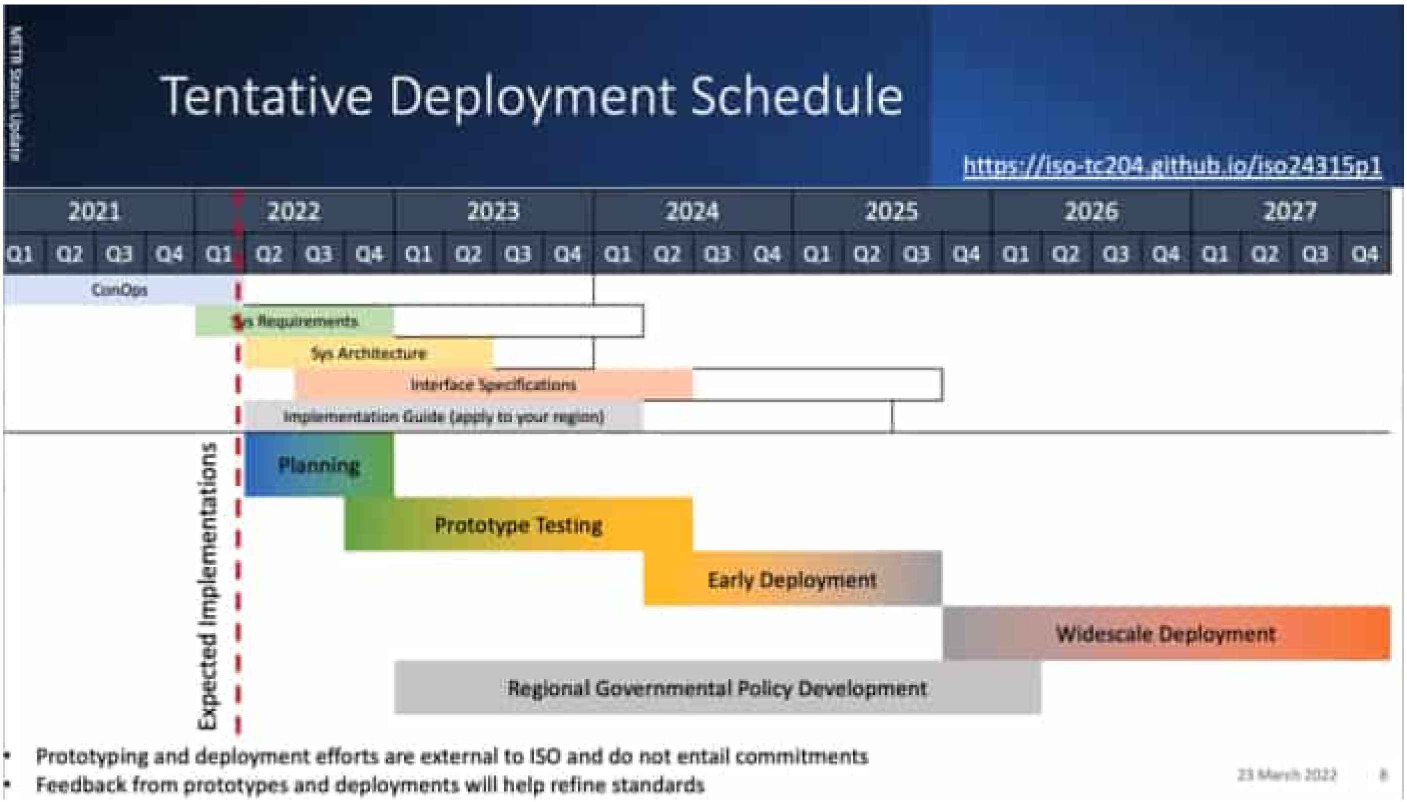

6

Vehicle User Requirements

6.1 At Issue

In addition to managing the steering wheel, accelerator and brakes, the human driver of a motor vehicle has several other responsibilities. These include:

1. Driver (or fallback driver) executive function—attention, alertness, sobriety, accident reporting
2. Passenger security—ensuring that passengers are and remain seated, belted, and within car seats (as appropriate)
3. Vehicle condition—ensuring the vehicle, with its components, is maintained, safe to operate, and that any load is properly secured

At present, these responsibilities are assigned exclusively to the driver.[22] This assignment requires re-examination in the light of ADS technologies. Increasingly, vehicles can "self-check" some of their components. Some vehicle conditions can be sensed and a decision about how to proceed in such cases can be automated. For example, sensors can tell if a passenger is belted, a data check could determine whether a vehicle is insured, and consistent responses to such situations, as defined in regulations, could be automated.

Conversely, there are some requirements that only a human can address, such as ensuring appropriate load securement, whether on the roof of a personal vehicle or in the cargo space of a commercial vehicle. In the absence of a human driver, it is unclear who will fulfil these requirements.

22 There is an exception; in the case of commercial vehicles, some responsibilities fall on the vehicle owner-operator. As an example of this, in April of 2022 a robotaxi operating in San Francisco was stopped by police because its headlights were not illuminated after dusk. In this case, the owner operator, GM, was considered responsible. https://arstechnica.com/cars/2022/04/cops-take-dim-view-of-autonomous-vehicle-driving-with-no-lights-at-night/

Finally, as ADS permits vehicles to become driverless, it is likely that human-only requirements will fall upon humans who are not present in, or proximate to, the vehicle but who exercise oversight remotely, i.e., a "teleoperator". Teleoperators will be at a distance from the vehicles they oversee, possibly even in a different jurisdiction. Telepresence requirements must be defined and enforced.

6.2 Policy Analysis and Recommendations

6.2.1 Requirements of Vehicle Occupants

We recommend that the regulator **identify, as part of any vehicle-user requirement**, **the entity responsible for meeting that requirement**, whether a human driver or corporate manufacturer, irrespective of any ADAS or ADS mediation, and prohibit the operation of a vehicle without an onboard responsible human agent until such time as this identification is complete.

In Chapter 4 on Responsibility, Liability, and Insurance, we recommended that the regulator remove any ambiguity regarding who is responsible for the *driving* of a vehicle. Here, we extend that recommendation to any requirement attached to the *operation* of a vehicle. Each vehicle user requirement that may be suspended for an ADS-equipped vehicle must be defined as such. Every other requirement must be explicitly assigned to the human agent responsible for that vehicle. That assignment must hold, irrespective of ADS capability or the presence of a human driving agent.

As an example, consider passenger securement. Vehicles often carry passengers who are unable to deploy a seatbelt personally, because of age, infirmity, disability, or state of consciousness. In such cases, these passengers must be secured by another before the vehicle begins to move. Today, this other is typically the driver. By defining this as the responsibility, not of a driver per se but of a human agent, the regulator will thereby be compelled to adjust enforcement protocols. Today, a traffic officer may ask a driver to secure themselves or their passenger. When there is no human agent present, it becomes unclear what that traffic officer is permitted and expected to do. In our view, making such determinations will be complicated, and the change difficult but necessary.

Another example of this is goods or property securement. In the event of a vehicle load, the same securement requirements as before should hold. Many jurisdictions carry out commercial vehicle inspections along highway routes or at border crossings. Regulations for any vehicle unaccompanied by a human must address the public need to carry out occasional inspection of goods securement, and the scope to address any defects that emerge. Here, again, we find a new regulatory complication. If an inspections officer is permitted to take action to secure a load, what level of training is required for that officer? What responsibility does the inspection agency now carry? If an inspector is not competent to address the load securement problem, what steps are then required of that inspector? What liability is defined for the owner-operator of a vehicle to redress any load securement failures? What rights does the inspection agency have to impound such a vehicle as a cost-recovery security? We raise these questions not to answer them, but to point out that the regulator's task will not be simple; assigning a responsible party, and then determining the consequences of that assignment, will require extensive effort.

In any case, we are certain that some degree of human oversight over vehicle operation will always be required, at least in the realistically-imaginable future. This might range from a fallback-ready driver, through a vehicle attendant, and perhaps ultimately to a teleoperator (about which more below). The regulator should be firm in assigning responsibilities to human agents where it remains appropriate to do so, irrespective of any pressure from manufacturers to relax these requirements.[23]

6.2.2 The Role of Self-Checking

Given our understanding of current and forecast automobile technology, we expect that vehicles will eventually be able to self-check the conditions of substantially all of their operating components. Such a vehicle will be able to reliably determine whether one of its sensors is broken or obscured, for example. Many of these self-checking capabilities may be product dependent or may not have applicable standards. They may not always be reliable, especially after repair, replacement, or modification. They may not be consistently present, or if present, may not perform to a standard on which a regulator can rely.

There will need to be some language in regulation that indicates a minimum amount of, and a minimum reliability for, system self-checking. Just as traditional vehicles (i.e., SAE Level 0) can advertise engine faults to a driver, so ADS-equipped vehicles must reliably self-check their sensors and software to prevent unobserved sensor or software failures. We therefore recommend that the regulator **define a minimum expected self-checking test regime**. This matter cannot be left to manufacturers to self-regulate, as they may be tempted to skew definitions toward what is suitable to their products, rather than what is safe for users. Regulators must insist upon, and indeed create and standardize, a common definition.

If any vehicle or vehicle subsystem is sold as self-checking, then a thoughtful guideline must be provided to understand what the vehicle user's responsibility is for that subsystem. This language would have to be generalized, so as not to be specific to a particular subsystem or type of subsystem. This may leave room for ambiguity which might generate problems for insurance settlement.

Many, but perhaps not all, vehicle conditions can be self-sensed and a resolution automated. For example, sensors can tell if a passenger is belted and a data check could determine if a vehicle is insured, and certain consistent actions, as defined in regulations, could be automated. While we admit that such processes lead to consistency and greater safety, we must also warn that they could lead to unintended consequences, including potential hardships for the passengers, such as an automatic seatbelt causing pain for a passenger whose arm is in a sling.

In the course of our research, we heard from experts whom we interviewed that regulators should draw a distinction between duties that pertain to the direct operation of the vehicle and those that do not. In the case of the former, we were told, regulators should consider assigning responsibility to an ADAS or ADS system to take positive, consistent action when they detect deviations from a standard. Examples of such matters include assessment of occupant safety and fitness to drive; checks that seatbelts are fastened; checks that the driver is not under the influence of substances that would impair judgment; and even that the driver is not in the grip of strong emotions. Such systems do exist or are being pioneered today, and typically involve the use of internal cameras to monitor the driver's state. As we have listed them, these systems begin with small infringements on the autonomy and privacy of the driver and end with more invasive ones. Regulators will be obliged to make a positive decision regarding the optimal trade-off between safe operation of vehicles and interference with the driver. Precisely where to set this trade-off will depend greatly on local cultural values; some polities place great stress on personal autonomy at the expense of collective safety, while others do the reverse. For this reason, we forbear making global recommendations even while asserting their critical necessity.

For duties that do not pertain to the direct operation of the vehicle, we were told by many interviewees that it is not reasonable to expect the vehicle to take any responsibility for them. No vehicle, for example, can ensure it has been loaded properly and securely, nor that it has been kept free of dangerous cargo (e.g., explosives, hazardous chemicals, etc.). These must remain the responsibility of a responsible human.

23 *A vehicle attendant or steward* would be a person working near or inside of a driverless vehicle to help passengers, load and unload goods, or to attend to load security. Such a person would not have any responsibility as a fallback driver but could carry responsibility for specified vehicle user requirements.

An option that we have considered, but do not recommend, is a requirement that an ADAS- or ADS-equipped vehicle become inoperable if its self-checking system indicates deficiencies. Such a requirement could prevent a vehicle's use in emergency situations. It is easy to imagine an impaired vehicle operating slowly, but safely, in a situation where it is impaired and where not using it could be more dangerous than using it: during a snowstorm or other adverse weather incident, in a situation where a licensed driver is injured but nonetheless must exit a dangerous location, or (for ADS) where no licensed driver is available but nonetheless it would be dangerous for the vehicle's occupants to remain where they are.

6.2.3 Requirements of Vehicle Teleoperators

We recommend that regulators begin immediately to **define a regime for teleoperation of ADS-equipped vehicles**, including rules regarding teleoperator-to-vehicle ratios, maximum time to secure attention, and so forth.

No matter how competent ADS-enablement becomes—even a century from now and even after the last fallback driver is removed from our vehicles—it will not be feasible to have machines roaming on roads and in cities without oversight. For this reason, we predict that as ADS becomes more sophisticated, teleoperation will become a standard feature of such systems. Such teleoperators could take over steering and braking under circumstances in which a vehicle leaves its ODD, such as the last few miles of a long-distance hauler as it crosses a border or enters a warehouse area.[24] They might be fleet operators who would be telepresent but not teleoperating during vehicle inspection, or for any case that requires a decision from an owner/operator, such as ensuring load security.

There is another, higher level of teleoperation at fleet level. Rather than being about steering or breaking, tending to the user requirements or recovery of an individual vehicle, it is about determining flow, congestion and interaction such as queueing among fleet members or across multiple fleets. For heavy good vehicles, it might involve congestion management, parking management, queueing at various nodes such as border crossings or inspection stations. For robotaxis, this implies the management of pick up and drop off as multiple providers queue for access in the same way that air traffic control systems manage runways. Systems like this are known as "orchestration" systems, and this is a critical area of work that is little discussed in the context of automated road vehicles. It will not be possible to optimize overall flow without addressing this, even if each individual vehicle is perfectly well-behaved.

24 We suggest this with uncertainty because the systems and methods for teleoperation and the human training and skillset to execute at least as safely as an in-vehicle fallback driver are not fully understood.

If we are correct and telepresence becomes commonplace, a solution presents itself to the problems described in a previous section, namely whom to designate as the responsible party for vehicle user requirements. But designating the teleoperator as the responsible party creates its own regulatory matters to solve—how attentive a teleoperator must be, how many vehicles a single teleoperator may monitor concurrently, how often teleoperation must sample the state of each vehicle, or whether teleoperation is, or can be, 'push' or 'pull'.[25]

These matters are complicated, and these requirements will change frequently as technology improves. The technology, and human requirements, of teleoperation is not yet sufficiently understood to write a teleoperations manual.[26] Nonetheless, we believe that any jurisdiction preparing for road vehicles operating under ADS should begin work on drafting a first edition of such a manual. Given the importance of teleoperation as a solution to so many problems facing automated driving—the majority of which are edge cases—the task cannot be delayed.

(The same argument applies to the management of last-mile delivery robots; see Chapter 14 for more.)

25 'Push' means that the human teleoperator watches and makes their own decisions about when to take over (this is currently how most sidewalk robots are teleoperated). It would be problematic for road vehicles to be reliant this way. 'Pull' means that the machine sends a signal "help me out here" and the teleoperator complies to assist or take over. This is how some sidewalk robots work, sometimes, and is better suited to teleoperating road vehicles than would be a 'push' approach. Realistically, teleoperation will always require both modes even though we would like to rely less and less on either, and especially less on 'push'.

26 Despite these difficulties, we suggest that, at least initially, teleoperators hold a driver's license of the appropriate class for any jurisdiction in which they are teleoperating a vehicle.

7

Training and Licensing

7.1 At Issue

Training is about equipping users with the knowledge and skills to safely operate motor vehicles on the road. Licensing is about ensuring who has had the requisite training, and as such can be relied upon to behave appropriately, or to accept enforcement of the rules. Both of these aspects are driven by, and are critical to, safety.

Today, we train and license users to directly operate motor vehicles with steering wheel, brake, and accelerator. This activity is fundamentally different from the activity of managing an ADAS or ADS, and surprisingly, the latter may be *more* difficult, not less. Managing an automated-driving system may involve maintaining readiness to yield to, or take over from, the system. It may also involve making decisions, or even responding to alarms, regarding where the burden of operation should lie. When that burden is with the human, the human must provide competent and timely operation.

Doing all this may require a higher level of cognitive executive function than the traditional dynamic driving task. The skill required is not currently fully understood, and that lack of understanding poses regulatory questions.

The training and licensing required of a user will differ among ADAS- and ADS-equipped vehicles. Just as there are similarities and differences between riding a horse and riding a camel, there are similarities and differences between driving an ADAS vehicle and being a fallback driver in an ADS-equipped vehicle. While the immediate problem we face is driver instruction and licensing for ADAS vehicles, the larger and longer-range issue to address will be registration and licensing for ADS-equipped vehicles. We first propose an ADAS approach, and then another, separately, for ADS.

7.2 Policy Analysis and Recommendations

7.2.1 Training and Licensing for ADAS

We recommend that regulators **amend their driver-licensing processes to include explicit instruction regarding ADAS and to emphasize that it is the driver's responsibility to understand any automated driving features they employ while operating a motor vehicle**.

There is a clear consensus in both the literature and among the stakeholders whom we surveyed and interviewed that existing training and licensing regimes are insufficient for these new technologies. ADAS is new, its capabilities are not widely understood, and marketing language has in some quarters bred complacency with this technology's power. Our interview partners were unanimous that no regulator should assume that existing training and licensing regimes are sufficient to regulate ADAS vehicles. It is necessary, we were told, that people understand what the sensors in their vehicle do, how they work, and how to maintain them.

When asked *when* this should be required, the response was emphatic: the time is now, for two reasons. Firstly, there is already a need for drivers of contemporary ADAS-equipped vehicles to understand the technology they are using and that to delay would engender risk. Our interviewees told us that the onus is on regulators, as the guarantors of public safety, to provide standardized training on the use of ADAS. Absent that training, it becomes more difficult to appropriately assign liability. ADAS users who misuse the technology, when called to account, could offer in good faith the defense that "nobody told me how to use it, it's unclear, and so that incident was not my fault". Such a gap undermines the accountability necessary to maintain safety and good order on the roads.

Secondly, the automated-driving industry is labouring under a burden of mistrust from users, many of whom have a weak, or no, grasp of what driving automation technology can and cannot do, and what threat it might pose. Providing training in ADAS today would help today's drivers prepare for the future, and build confidence in the technology now, when the industry needs it most.

We underscore these findings and believe that regulators must begin now to ensure that future drivers, before receiving their license, come to grips with the promise and pitfalls that ADAS offers.

Unfortunately, our certainty must confront the fact that there are myriad ADAS-equipped vehicles, from a variety of OEMs, and with more variety likely to come. No standard driver-training-and-licensing process could categorize them all, nor could any regulator reasonably update its training in a timely fashion. Finally, it would be burdensome for drivers to receive training on so many possible ADAS, especially since a vast plurality (if not majority) may never drive a vehicle so equipped.

In keeping with the principle that we articulated in Chapter 4.2, in making our recommendation we choose not to relieve the vehicle consumer of responsibility. Those consumers are the parties best equipped to educate themselves on the capabilities of the vehicles they drive—whether purchased, rented, shared, or leased—so the onus is upon them to do so, not on the regulator.[27] That said, it is reasonable to suppose that many vehicle consumers will not do so, thereby blurring the lines of accountability away from the consumers and onto OEMs. The best solution to this problem, in our view, is for the regulator to require, as part of its licensing process, an explicit discussion of ADAS in the abstract. The point of this discussion should be that while ADAS can be quite helpful, ultimately it does not absolve the driver of responsibility or liability. For that reason, any driver must familiarize themselves with their vehicle's capabilities (or lack thereof) before engaging any ADAS feature. In this way, we hope to ensure that road users are equipped to behave knowledgeably and responsibly. (We add that our recommendation in Subsection 4.3—that action be taken against misleading marketing as to the extent of an ADAS-equipped vehicle's capabilities—will also help in this regard.)

We considered recommending that regulators insist that no vehicle be permitted on the roadways that require additional training; i.e., that any ADAS system that is sufficiently complex as to require instruction in its use is not suitable for use on the roads. We rejected this view, as we think this would be infeasible until the highest levels of driverless capability are reached, and non-automated vehicles removed from the road. Such a time is decades away, if not even further off, and we think that ADAS has too much potential to reduce harm for it to be restricted in this fashion.

7.2.2 Training and Licensing for ADS

At present, the training and licensing regime described above is for humans only. The implication of ADS technology is that such a regime is insufficient. Accordingly, we recommend that regulators **begin to design a consistent licensing and registration process specifically for ADS**.

Today, in most jurisdictions, a vehicle travelling the roads has two components with respect to its legal operating status. Firstly, it is a registered vehicle with properties of size, make, model, and unique identification (typically a Vehicle Identification Number, or VIN).[28] Secondly, it has a driver, who holds a specific class of driver's license for specific conditions of operation; a license may be suitable for operation of an automobile but not a motorcycle, or for a vehicle carrying no more than ten passengers, and so forth. The registration status of a vehicle and the license status of the driver are independent.

A typical family vehicle will have a single registration, but it may also be operated by two or more licensed drivers. Each of those drivers holds independent licenses and those licenses may be of different classes. A traffic violation is typically associated with a licensed driver. A key reason that vehicles in moving violation are apprehended is to identify the driver. Issuing citations without vehicle apprehension (a matter we will discuss in more detail in Chapter 8) means the citation would be attached to the person to whom the vehicle plate is registered, rather than the driver.

An ADS-equipped vehicle has a third component requiring legal identification, namely its ADS. That ADS, when engaged, is the driver of the vehicle. The act of dis/engaging an ADS is theoretically similar to two licensed drivers switching places in the driver's seat. This is not the same as engaging an ADAS, which is analogous to a single licensed driver engaging cruise control.

An ADS, then, is an independent component. Just as a physical, human-operated vehicle has a make and model, and when travelling has a driver, an ADS also has properties conceptually similar to each of these (even a driver, if dis/engageable). These properties include an ID of its ADS class (ADS-H, ADS-V, and ADS-DV), and an ID for its most recent software update, which in turn would include an ID for its provider. It is possible that other descriptors will be required.

In this way, an ADS-equipped vehicle that requires or permits a fallback driver (ADS-H, ADS-V) will require three registrations: one each for vehicle, driver, and the ADS. An ADS-equipped vehicle that permits no fallback driver (ADS-DV) will also require three registrations: one each for vehicle, fleet owner, and the ADS.

A new and qualitatively different requirement arises because of the impact of software updates. Such updates may change the capability, and possibly the liability profile, of an ADS. It is also possible to imagine an upgrade from one ADS class to another which would automatically have significant impacts on the liability profile of that vehicle-driver-ADS or vehicle-owner-ADS triplet. This comprises further motivation for the critical definitions of the liability switching boundary between human and ADS discussed in Chapter 4.

Hence, regulators will require an additional licensing and registration regime for ADS-equipped vehicles. It will be possible for this regime to impose requirements in real time, unlike today's static processes for vehicles or human drivers. Because of their implications for liability, such updates must be accurate and timely, and regulators will be obliged to promulgate rules that ensure that both manufacturers and users make, and receive, updates reliably and quickly.

27 We comment further on driver awareness of vehicle capabilities in Chapter 4.

28 VINs are defined by the International Organization for Standardization in ISO 3779 (content and structure) and ISO 4030 (location and attachment).

8

Enforcement and Emergency Response

The complexity of traffic enforcement and emergency response is likely to grow as ADS-enabled vehicles become common. Notwithstanding this increase in complexity, we expect that the direction of both enforcement and emergency responses over the next several decades, and after a disruptive period of mixed traffic, will tend toward a lower volume of enforcement and emergency responses. In this chapter, we offer some recommendations to ease this transition.

For the regulation of enforcement, we propose several new capabilities. Some of these may be exercised remotely, such that officers, or roadside systems, may detect infractions, gather evidence, and then deliver citations to the vehicle owner electronically, without necessarily interacting with a human. Others may involve interaction with a human teleoperator.

For emergency response, we observe that, absent a responsible human at the scene, ADS will require changes in EMS responder protocol. New vehicle designs may necessitate new vehicle entry tools or training, especially considering the disposition of passengers in altered seating and restraint designs; new materials and fuels may require new protocols for dealing with hazardous environments.

(One area that we do not explore, but mention in passing, is the need to consider emergency enforcement in the case of driverless vehicles commandeered for criminal purposes. We cannot fully anticipate how this phenomenon will feature in the roadways of the future, and consequently we are unable to suggest solutions. Nonetheless, this is an issue that requires scrutiny.)

8.1 At Issue

Enforcement of traffic regulations helps to ensure safety for all road users. Emergency response to incidents, such as crashes, aims to:

- Reduce the harm outcome of such incidents;
- Gather data to be used to enhance safety (i.e., to consider what safety-related matters may have been learned from the incident); and
- Help identify liability.

The advent of ADS will affect these matters as follows.

Enforcement procedures: How are driverless vehicles to be halted, opened, entered and, if required, impounded? There will be cases with and without a human present within or on the vehicle. There may be cases where a vehicle is not in communication with a teleoperator that can help with procedural compliance (communication failure), or with a teleoperator whose actions are able to defeat the officers' attempts to control the situation or vehicle. There may be cases where the ADS is experiencing an error or is under the influence of a bad actor.

Officer training and emergency procedures: It appears to be the case that officers attending crash sites are beginning to encounter vehicle configurations and systems that may defeat current methods and procedures to enter the vehicle, extract passengers, extinguish fires, etc. Will new vehicle designs increasingly challenge the training and procedures of attending officers? Will new guidelines, procedures, and training programs be required? If so, what regulations and policies are needed to define and disseminate such programs?

Vehicle ID and emergency procedures: It appears that the ability to identify a vehicle can be hampered by the loss of vehicle plates during severe crashes. Vehicle ID may become more critical for the officers who are attending a crash to pre-determine the nature of the vehicle subsystems so that they might apply the appropriate methods to perform their task. Some of this information will be required in real time to inform emergency response.

Harmonization: Will this lead to cross-jurisdictional issues? It might, if an out-of-jurisdiction vehicle is not registered in the appropriate emergency response system. It might be necessary to develop a physical, harmonized vehicle ID system that would be readable from a distance and at speed (as is a registration plate), and instantly associated with sufficient descriptive information that would identify the party liable for violations, the level of ADS installed in the vehicle, the teleoperator connection, and any vehicle properties related to fire, chemical hazard, or emergency evacuation. This would require a harmonized plate system that would span international boundaries and readers connected to a method to provide the required information for enforcement officers within only a few seconds. While vehicle automation might reduce the incidence and severity of some crashes, it may complicate crash preparedness and response.

8.2 Policy Analysis and Recommendations

8.2.1 Enforcement

For both ADAS and ADS-H-equipped vehicles, existing enforcement procedures will satisfy; innovation is not required, as these must feature a responsive, on-board fallback driver. For ADS-DV vehicles (i.e., ADS-equipped vehicles that have no driver or fallback driver, as described in Chapter 3.2), and ADS-V vehicles within their ODD, the existing regime will be insufficient.[29]

For that reason, we recommend that the regulator **adopt law-enforcement interaction protocols specifically for vehicles under ADS-DV and ADS-V operation**. These protocols should stipulate that:

1. **ADS-DV and ADS-V vehicles must have a unique external marking that can be read and interpreted at maximum roadway speeds for remote identification**

An obvious solution to this problem would be to provide visually distinct license plates for such vehicles, but this may not be the best approach; see Section 8.2.2 following.

2. **Enforcement officers or systems are empowered, upon vehicle identification, to issue electronic citations to ADS-DV and ADS-V vehicles without apprehending the vehicle**

On their own, these recommendations will assist enforcement personnel to address routine infractions, such as vehicles operating with broken lights, exceeding speed limits, and so forth. It may be the case that, as ADS technology reaches maturity, driving systems will be sufficiently sophisticated as to always obey speed limits, or to monitor their own lights and signals and switch to "safe mode", or even pull over, when these are inoperable. In such cases, the number of such infractions and citations may dwindle over time.

While the sophistication of this technology will increase, the possibilities of malfunction, misuse by bad actors, or use in criminal enterprise will always exist. For those reasons, we recommend:

3. **ADS-DV and ADS-V vehicles must have a teleoperator that enforcement personnel may** contact in real time for the purpose of having the vehicle exit the roadway, come to a complete stop, and permit an enforcement officer to open the vehicle for inspection and apprehension of any passengers
4. **Teleoperators must comply with enforcement officers**, just as human drivers of non-automated vehicles are so compelled

For instances where a vehicle's connection to a teleoperator has failed, the teleoperator is a bad actor, or a teleoperator's operation of a vehicle is compromised, we recommend:

5. **Enforcement officers or systems must have access to an interdiction system by which they may cause ADS-DV and ADS-V vehicles to safely come to a stop and disable themselves**

In the course of our research, one expert whom we interviewed dismissed as "science fiction" the idea that law enforcement could have the power to directly disable an automated vehicle. Notwithstanding that impatience, we believe that contacting a teleoperator for this purpose, and ultimately to have access to an override, is a reasonable expectation, both of expected technology and, to the extent that we understand it, within current law.[30]

These recommendations are consistent with current practice, at least in some jurisdictions. At time of writing, each of Arizona, New York, and California require that AV manufacturers prepare a "law enforcement interaction protocol" that will instruct first-responder personnel how to interact safely with the AV. California goes further and specifies that these protocols determine how to recognize a vehicle operating in ADS mode, to safely disengage it and know that it has been disengaged, and to communicate with a teleoperator; and to make these protocols available via the Internet to facilitate fast and simple access by law enforcement. We recommend that the regulator **adopt these or similar protocols for ADS-DV and ADS-V vehicles**.

We also recommend that **such protocols be harmonized to every extent possible**. Consider a scenario in which an ADS-DV- or ADS-V-equipped vehicle is licensed in and teleoperated from one jurisdiction (say, Washington State in the USA) but becomes subject to a law enforcement interaction in another (say, British Columbia in Canada). It would be required for the teleoperator in Washington to understand the protocol for British Columbia, or it would be necessary for a protocol

29 Note the near-equivalency of an ADS-DV vehicle and an ADS-V vehicle in its ODD. The only material difference is that the ADS-DV vehicle may not have any driver aboard but the ADS-V vehicle would have an alert licensed driver aboard in all but exceptional circumstances (e.g., medical emergency).

30 On April 10, 2022 a robotaxi with no in-vehicle fallback driver (ADS-DV, by definition) was stopped by police in San Francisco because its headlights had malfunctioned. The vehicle "yielded to police and then pulled over to the 'nearest safe location.' One of the officers contacted [the teleoperator] after the traffic stop," and the matter was rectified. What is telling in the referenced report is that the teleoperator "directed the car to pull over at [a nearby] location — across the intersection — when it became clear that the car was the subject of a traffic stop **and the officer was clear of the car**. When asked, the spokesperson declined to reveal whether [its] vehicles would behave differently if the stop happened on a highway versus a city street." (Wiggers, 2022; emphasis added) Our concern is that regulations and officer training for such interactions be established. It must be certain that any apprehended vehicle not move (whether by error or mischief) after it has brought itself to a safe location. Hence, while under a traffic stop, the ability of the vehicle to move must be strictly determined by an enforcement protocol.

to switch teleoperation centers on crossing the relevant border. The latter case may seem more desirable; our (tentative) opinion is that such a requirement would require a technological protocol that would pose a high risk of abuse and exploitation by bad actors.

Beyond this point, our recommendations must halt, as we are not qualified to address the many legal questions that arise. Rather than recommend measures, we will simply raise matters that will require legal expertise to determine. These matters include, but are not limited, as follows:

- ADS-DV vehicles will often not have a licensed driver on board and may indeed lack passengers entirely; in such situations, were an officer to issue a citation for improper operation, there may be no person competent to challenge it
- Enforcement officers must have a way of being certain that stopped ADS-equipped vehicles remain stopped, given the potential for remote operators to weaponize a vehicle
- While it is imperative that teleoperators of ADS-equipped vehicles be required to bring a vehicle to a stop if an enforcement officer so requests, the scope of a teleoperator's obligations beyond this point must be determined
- Enforcement officers will require guidance, and likely explicit rules, in respect to searching a vehicle that may not have a proximate human observer of that search
    - Consider the circumstance in which a passenger vehicle, without passengers, is suspected of transporting contraband. Because of the potential opportunities to use ADS mobility technology for illegal purposes or illicit gain, and for the potential of illegal activities on the part of enforcement officers, combined with any unknown complicit connection with the teleoperator, there are several complex scenarios that regulators must define and then establish protocols for their resolution

In the course of our research, we found that experts and stakeholders with whom we spoke were adamant that law enforcement receive, as a matter of course, access to the data generated by ADAS and ADS, to ensure that law enforcement personnel, when investigating road incidents, are able to quickly and unambiguously determine if a human operator's claim that the vehicle's systems, rather than human behaviour, are at fault. Public representatives must be able to discern who or what was operating the vehicle during a crash or a violation. Both manufacturers and regulators must define, standardize, and certify methods to determine and record dis/engagement events. Beyond this, we are silent; the questions of when and how this data is to be made available to law enforcement are important, but we lack the expertise to settle them. For more on this difficulty, see Chapter 12 on data collection.

There are likely many additional scenarios that regulators should enumerate and resolve before ADS-DV or ADS-V vehicles begin operating on public roads.

8.2.2 Emergency Response

We recommend that regulators **require manufacturers of ADS-DV vehicles determine rules and procedures for responding to an emergency affecting ADS-DV vehicles** where physical presence of an emergency crew at the site of the emergency is required and make these available over the Internet.

ADS-DV vehicles may be disabled or involved in a crash that requires an emergency response at the roadside. It is critical for the emergency crew that is addressing an incident to be able to fully understand any hazards interacting with such a vehicle entail, such as special fuels, battery fire risks, unusual entry requirements (Jaws of Life), risks of fire that cannot be extinguished using conventional means, and others. There must be a way to inform the emergency crew of such risks within a few seconds of recognizing an ADS-DV vehicle.

The most straightforward solution, building on recommendations in the previous chapter, is that, firstly, ADS-vehicles feature a unique external identifier that can be read and interpreted upon arrival of emergency personnel. We were told during one of our interviews that severe crashes can dislodge license plates, throwing them away from the vehicle and making them unlocatable. The interview subject suggested that such vehicles should have, in addition to (or in lieu of) a distinct license plate, a marking, label, or plaque that cannot be easily dislodged. Secondly, that manufacturers prepare guides to the risks and hazards their vehicles may pose, and publish these, so that first responders can quickly and easily understand the situations they may encounter.

Against this, some have suggested that clearly legible markers advertising a vehicle's status as an ADS will prompt abuse by other road users, who will take advantage of the system and its commitment to safe operation to seize the right of way or otherwise gain advantage. If so, this argument supports the creation of ADS-only spaces, and not foregoing the clear advertisement of a vehicle's ADS capability. Safe and legal operation, in our view, may require such a commitment.

9

Aftermarket Modifications

9.1 At Issue

Relative to traditional ("Level 0" or "Level 1") vehicles, automated vehicles feature significantly more software and specialized sensors, a ratio that will only increase in the future. Because of the integration between these software and hardware components, aftermarket modifications that directly or indirectly impact software, sensors, or the interaction between them might easily have unintended consequences.

9.2 Policy Analysis and Recommendations

We recommend that regulators **restrict aftermarket modification** that could alter any aspect of an ADS **to be made only by the entities that will be held liable for the vehicle when its ADS is engaged**. This must be the case in order to align with our recommendation that the ADS provider be liable when the ADS is fully engaged. It is not reasonable to ask an ADS provider to be responsible if there has been a system modification that influences the efficacy of their ADS.

We do not make this recommendation lightly. We are well aware that such a regulation shifts substantial market power away from consumers to OEMs, who will enjoy monopoly rights over aftermarket modifications and repairs, and will, we presume, take advantage of that position to set prices accordingly. We are also aware that this recommendation is at odds with the right-to-repair philosophy. Nonetheless, we do not see a reasonable alternative to this approach.

We take this position because the consequences of a different course of action are too much for a regulator to bear. For any vehicle with an ADS (ADS-H, ADS-V, or ADS-DV), it must be the case that all parties—owners, passengers, enforcement officials, regulators, proximate vehicles and pedestrians, and more—must be entitled to assume that the ADS will perform as expected and that this assumption overrides the right-to-repair. Absent such an assumption, any ADS-equipped vehicle would pose a burden of wariness and care on everyone that interacted with it while in operation. Such a burden would be insupportable.

As a consequence, then, no modification can be made that could compromise the expected performance of the ADS. Should such modifications occur, and the liable party or parties be able to determine that such modifications have been made, the liability and responsibility allocation for the ADS would be compromised, and the insurance subrogation process disturbed. It is clear to us that the public as a whole, and the regulator in particular, have a strong interest in bright-line allocation of responsibility, and that interest trumps other concerns.

We note that some regulators have adopted positions like ours. The Australian government, for example, has taken the view that the greatest risk in this domain comes from private individuals installing aftermarket ADS kits themselves, who would be unlikely to have the knowledge to identify and adequately address safety risks such as managing cybersecurity risks or failure of the ADS to function properly by regularly updating software. Therefore, that regulator prohibits individuals from installing aftermarket kits to vehicles, will introduce a new legal offense of "third-party interference to an ADS", and will require installers to be accredited, which will include meeting corporate obligations, such as a corporate presence in Australia.

In Chapter 4: Liability, Responsibility, and Insurance; we discussed the importance of the secure performance of an ADS to the reasonableness of holding the ADS provider liable.

For the reader's convenience, our recommendations in this regard included the following:

1. No seller, lessor, renter, or vehicle maintenance operator may modify any part of an ADS, including any of its software, sensors, effectors, connectors, or connections
2. In the event of a repair, alteration or recall of the ADS components of an ADS-equipped vehicle, only a maintenance provider approved by the party or parties that are liability for the ADS when the ADS is engaged can make the required adjustments
3. Any prohibited tampering, misleading marketing, prohibited modification, repair, or alteration to any ADS aspects of a vehicle must be regulated as a criminal offence

There is scope here to advance the right-to-repair to a certain extent. Regarding point 2 above, we wish to emphasize that while ADS manufacturers are responsible for modification or repair of the system, they may also delegate this responsibility to approved agents. The regulator might insist that manufacturers of ADS-equipped vehicles that have been manufactured for private ownership be required to consider delegating this agency to qualified parties who seek it and can demonstrate their fitness for it, and that approval of that delegation not be unreasonably withheld. In such cases, the market for aftermarket repair will be less of a monopoly, tilting the balance of power away, to some extent, from OEMs and toward consumers.

We are aware that simple maintenance activities such as repairing dents and scratches or using a mechanical car wash on an ADS-equipped vehicle might alter its sensors. An Internet search for the terms <autonomous vehicles car wash> provides evidence for this risk. We do not have sufficient expertise to enumerate which features could be modified in the aftermarket. But again, we think that a balance must be struck.

For example, one might judge that the owner of such a vehicle ought to be able to have it painted any desired colour. However, colour impacts reflected light, and reflected light may change the properties of the image sensors. Knowing this, an unintended consequence of our recommendation could be that OEMs will expand their definitions of the scope of the ADS-related aspects beyond reason, to make more elements of the vehicle exempt from aftermarket modification and free-market competition for the same.

We recommend that **the provider of the ADS must state which ADS aspects are to be prohibited from modification**; that the regulator **take appropriate action to prevent such modification, subject to a test for reasonableness**, to ensure that the power of this regulation is only deployed on behalf of system elements that truly affect ADS performance.

10

Automated Trucking and Cooperative Truck Platooning

While much of the discussion of automated vehicles (AV) has related to its effects on personal travel, by some accounts AV's impacts on commercial delivery of freight are likely to arrive sooner and have more dramatic effects on the conduct of business in the logistics sector. This subject clearly deserves significant attention.

Here we consider the impact of automation on trucking in regard to two issues: the automation of Heavy Goods Vehicles (HGV) and the electronic tethering of those vehicles into short road trains, known as *cooperative platooning*. These two matters are related, with the first being an important enabler for the second. A form of the latter, referred to as Long Combination Vehicles or LCVs, is done now by physically tethering non-automated vehicles.

The safety, regulatory, and inspection issues for cooperative truck platooning will almost certainly be different than those for LCVs. For the present, we recommend focusing only on the simplest cases of highway platooning.

10.1 At Issue

The spectrum of issues regarding driving automation for HGV differ considerably from those for human passenger applications. Most of these differences make it more likely that ADS for trucks and cooperative platooning will become pervasive before ADS for human passenger vehicles does. These differences include, but are not limited to, the following:

1. Since it is not marketed to the public, automated trucking is subject to far less consumer hype, which in turn means the issues surrounding its use are relatively better understood and discussion of them is less fraught
2. The ODD for commercial truck ADS and platooning is dominated by divided highways which are easier and safer to navigate using ADS than urban and suburban streets. There will be fewer 'edge cases'
3. Land use economics, and allied zoning rules, have impelled the creation of most contemporary logistics and warehousing areas convenient to such highways. It is conceivable that automated trucks might operate exclusively on a subset of the road network essentially characterized by logistics areas and highways. Such a development might permit the sooner arrival of commercially -viable AV trucks, as the operating environment would tend to be small, applied to comparatively few highway entrance areas and offramps; such a development could offer massive commercial benefits to logistics companies, which are consequently motivated to hasten its arrival, especially compared to the more distributed and continuously-delayed social impacts of passenger vehicles
4. The state of cooperative-platooning technology suggests that participating companies have the potential for quick, measurable impacts in operating costs (i.e., more value per vehicle-kilometre) even without going completely driverless, a situation that will not immediately apply to passenger ADS. The prospect of great economic benefit also encourages market participants to spur the technology's arrival
5. Over the long term, platoons would likely permit reduction of labour costs at scale, particularly given the widespread, and widely acknowledged, labour shortage in trucking and the logistics supply chain, noted in Australia, Canada, Europe, the United States, and other jurisdictions[31]
6. Energy savings at scale are also likely to be significant; controlled tests of platooning using ICE vehicles have indicated fuel savings of between 4% and 10%, while some

31 https://www.yourtravis.com/news/truck-driver-shortages-in-europe-and-the-us-and-what-countries-and-companies-do-about-it/

commercial claims have been as high as 15%. One Tesla spokesperson claimed a cost drop to $0.85 from $1.26 per mile for Tesla's electric trucks, a 32% improvement[32]

7. The flexibility of operations for HGVs (particularly at night, which is currently constrained by human ability) could result in reduced congestion, more efficient use of existing road infrastructure, and improved freight operation times

Taken as a whole, the economic pressure to optimize goods movement is greater and more focused than the economic pressure to optimize passenger movement. Put another way, the requirement to optimize goods movement is critical and immediate, while the requirement to optimize passenger movement is complicated and slowed by the debates regarding private vehicle ownership, shared vehicle use, public transportation, and active transportation, none of which applies to the optimization of goods movement. For these reasons, coupled with consistent delays in the readiness of driverless passenger vehicles, automated trucking and especially cooperative platooning have recently become far more attractive relative to passenger vehicles as near-term opportunities for investment and deployment. This market perception is likely to continue.

This attraction is tempered by the fact that HGVs are on average an order of magnitude heavier than the average family vehicle, making the potential kinetic damage and harm of an incident involving ADS-equipped HGVs dramatically greater than a parallel incident involving ADS-equipped passenger vehicles. Thus, regulators of this sector must observe the maximum possible concern for operational safety.

One item we heard consistently from our interview subjects is that the public is fascinated by AV technology, but also afraid of it. Nowhere does this fascination and fear apply more than in the application of ADS technology to trucking. People outside of the sector are already nervous around big rigs, and if automated HGV platoons are permitted, that anxiety would likely be compounded. The biggest question that the industry faces—and by extension, governments—is whether the public trusts driverless trucking to be safe. If it does not, then much heavier regulation will be required, and consequently slower uptake of the technology and foregone benefits.

Therefore, we were told that what government and industry needs to do now is to build trust. Accordingly, governments should consider a variety of regulations, with an expectation that these can be revisited in the future and relaxed when the public is ready for that to happen. These might include, but are not limited to, speed regulators on AV truck platoons, restrictions on lanes where platoons may operate, caps on the numbers of vehicles that may participate, and so forth.

Conversely, governments should not hesitate to begin permitting safe experimentation. Pilot projects to advance platooning and other AV trucking technologies are important given the ongoing driver shortage. A strong commitment to safety is required, but if this commitment is being respected, innovation in this sector should be permitted, and even encouraged.

32 Roberts, Jack, "Truck Platoons on the Horizon?", www.truckinginfo.com, 2018

10.2 Policy Analysis and Recommendations

The stakes of pervasive deployment of cooperative platooning are particularly high. This implies that the urgency for the regulator to address policy and regulations for such vehicles will likely be greater and sooner than the urgency to regulate ADS for passenger vehicles. Hence, we recommend that regulators **prioritize their regulation of ADS deployment to permit ADS for trucks and cooperative platooning**. We further recommend that regulators **permit such deployment in well-understood and tightly-controlled stages**.

The technologies for ADS-trucks and cooperative platooning are quite different, and must be treated separately.

10.2.1 ADS for Trucks

We expect that ADS-H or ADS-V technology will be deployed in trucks long before ADS-DV will be, meaning that human drivers will still be required. These systems will reduce human effort and make operation safer, but only if, when acting as the fallback driver, human drivers are always responsive when the ADS requests that they take over the dynamic driving task (DDT). Thus, the central regulatory problem with ADS for trucks is whether the attention of a human fallback driver is always available. We recommend that regulators **insist upon appropriate measures to ensure a timely and consistent human response to ADS disengagement and verify that this requirement is respected**. These measures must be negotiated with industry representatives and might include in-vehicle sensors to monitor engagement levels, remote checks by a teleoperator, or other approaches.

We recommend a **conservative program of registrations for ADS deployment**. For example, a jurisdiction might begin by permitting the use of this technology to a few fleets registered to a small number of operators, who would need to achieve a pre-agreed level of performance with regards to safety. If these targets were met, operators might register a larger fleet, again to achieve a sustained history of safety but to a higher standard, and so on in this fashion. In other words, fleet operators deploy ADS only as fast as the fleets can unequivocally prove themselves to be operationally safe at each stage.

Unfortunately, such fleets will not be able to prove that its ADS is operating in a sufficiently safe manner simply by driving many miles with a small number of unplanned disengagements. For this reason, we recommend that certifying authorities for ADS and for truck platooning **acknowledge that it is inappropriate and impractical to prove road safety for ADS via road-testing alone**. As noted by Dr. Philip Koopman, it is unlikely that any regulator could accept a driving system with or without a human back up driver that is not "at least as safe as" a competent human driver. How safe are such drivers? As measured by fatality rates, human drivers in the USA are responsible for approximately one fatal crash per hundred million miles, as per Koopman. Canadian drivers are slightly safer, being responsible for 0.8 fatalities per hundred million miles. Thus, according to Koopman, to have 99% confidence that a driving system is at least as safe as an American human driver, we would require more than 460M miles of driving without a fatal crash or 664M miles with one crash; and even higher to meet the Canadian standard.[33]

To compound this, all these test miles must be driven using the *same* system, namely software, hardware and ODD, that an authority plans to certify. Clearly, delivering so many test miles for a single system is untenable on physical roadways, much less for multiple iterations.

The solution is a combination of extensive simulation, aligned with a significant number of miles on representative ODDs (roadways and conditions). We add, *pace* Koopman, that the supporting digital miles must not only be representative of the intended ODD miles but also provide the opportunity for what Koopman calls "interesting miles" in which all the "rare situations that are seldom seen" can be tested.[34] Hence, before any program of certification is undertaken, the simulation systems to be used must themselves be confirmed, a task that will require significant specialized knowledge.

We recommended that initially, **existing rules for the number of hours of service that human drivers are permitted to drive in a given period (HOS) should not change**.[35] This may seem paradoxical, given that the point of automating the DDT is to use fleet vehicles and their drivers more efficiently. While we agree that this is the ultimate goal, we are in the early days of this technology, and the effort required of human drivers to remain fully attentive while only watching the road, rather than driving on it, will be challenging. This challenge will be sufficiently great such that even operating to the extent allowed by current HOS while remaining free of incident will require effort. As the technology matures, as well as the skill of backup drivers at remaining unstressed while also ready to take command, HOS rules might be relaxed, but this relaxation must not occur before this maturation and skill have been conclusively demonstrated. Such a demonstration may allow for evolution in HOS rules. An example of possible evolution would be that the regulator permits a driver, who has driven some large number of kilometres without any irregularities in dis/engagement, the privilege of working longer hours. Depending on the jurisdiction, this privilege could be expressed in several fashions, such as the award of advanced levels of commercial driving licenses.

Eventually, ADS-DV will be deployed to trucks. We are confident that this level of driving automation will one day be achieved, even though it may be limited to vehicles that operate predominantly on divided highways. When this is achieved, a human driver will no longer be needed for such vehicles. This does not mean human attendants will be entirely superfluous; humans would likely accompany or rendezvous with such vehicles for matters of load security, load inspection, and border crossing, except where a teleoperator may suffice.

33 Koopman, P. (2022) How Safe Is Safe Enough? Measuring and Predicting Autonomous Vehicle Safety.

34 Knowing that enough edge cases have been found and incorporated into the simulation system will be itself a difficult task, as such cases seem to multiply without limit.

35 HOS is an acronym for Hours of Service, the term used for these rules in Australia, Canada, and the United States. In Europe the term used is Driver's Working Hours.

We recommend that the regulator **not permit truck ADS-DV deployment until a pre-determined proof of safety program for ADS-H or ADS-V driving has been completed**. This recommendation maximizes public safety while aligning the incentives of OEMs to take safety seriously as a principal goal of automating their fleets. For more on how a government might log and monitor such an outcome, see Chapter 12 on data.

10.2.2 Cooperative Platooning

Platooning of HGVs, if it is judged ready to be broadly deployed, may provide economic value in the form of safety, fuel savings, reduced stress, and labour benefits.

We recognize the promise of simple ADAS level 2 truck platoons where a follower vehicle is electronically locked into the lateral and longitudinal controls of the lead vehicle. Jurisdictions studying such deployments currently require a fallback driver in the second (or third) vehicle, but it is reasonable to contemplate a driverless (ADS-V or ADS-DV) follower vehicle, and ultimately multiple follower vehicles, as this technology matures. We believe that constrained versions of this approach will survive scrutiny for safety and efficiency in the foreseeable future. For that reason, in the following sections, we offer recommendations regarding longer platoons, mixed platoons, cross-jurisdictional platoons, platoons with only a lead driver, and more. These are more complex, and far riskier enterprises than simpler ADAS Level 2-led platoons.

10.2.3 The Range of Platooning Regulations

In our view, the regulator should constrain cooperative platooning to only the simplest configuration in the near term and **regulate more complexity in measured steps by ensuring extremely reliable levels of safety** before proceeding to further stages, which might include longer trains, mixed vehicles, removal of a driver in follower vehicles, and so forth.

Conceptually, cooperative truck platooning ranges from the very simple to the very complex. In its simplest and least risky form, a platoon features two vehicles from a single company/operator each with an attentive driver responsible for their own vehicle, travelling within a single jurisdiction, insured under a single policy, teleoperated by a single management center, with both trucks' ADAS integrated and updated from a single provider. This case is the most straightforward to regulate. For that reason, regulators should initially permit this alone; only when this use case has been mastered will it be appropriate to attempt more.[36]

At the other extreme, the very complex case would be an arbitrary number of vehicles (i.e., more than three), from a plurality of companies, only some of which have an onboard driver (for example, six vehicles and three drivers), vehicles registered in various jurisdictions, insured by multiple insurers (implying a very complex subrogation process), each operated by an independent teleoperation center, and each supplied with ADAS or ADS from an arbitrary provider.

36 https://guides.loc.gov/trucking-industry/autonomous-trucking

Such a complex case may never occur, and if it does, it will almost certainly not emerge for decades in the future. However, manageable configurations between these two extremes may occur sooner. For that reason, regulators should develop constraints for platoon length greater than two, multiple companies, multiple jurisdictions, multiple insurers, multiple teleoperators and multiple ADAS/ADS providers.[37]

Many matters relating to platooning are, in our view, purely local, including inspection, or the collection, frequency, and retention of data relating to platoons. Some matters, though, will require harmonization, notably including forming platoons of mixed-registration vehicles. Others will require private-sector firms to reach agreement among themselves, with the public sector remaining neutral on the solution but insistent that one must be proven safe before platooning may begin. These include providing for the interoperability of competing makes and models of truck and platooning ADS.

The principal regulatory question that a jurisdiction must settle is to what extent a platoon of a given number of vehicles may consist of homogeneous or heterogeneous entities and operators. Complicating factors include:

1. A platoon may be comprised of vehicles registered in multiple jurisdictions or from different operators
2. Liability must be apportioned
3. Insurance subrogation must be clarified
4. Tolls must be assessed and collected
5. Vehicle weight rules must be applied
6. Weigh-in-Motion systems must be adjusted to account for assignment of costs or fines

We recommend that **platooning be harmonized across jurisdictional boundaries to the greatest degree possible**. Having an undue number of jurisdictional constraints on the composition of platoons and the continuity of their trips across jurisdictional boundaries will diminish the operational and optimization values of this technology, especially if longer platoons are contemplated.

10.2.4 Be Cautious of Long Train Cooperative Platoons

Initially, at least, we recommend that regulators **limit platoons to two vehicles sharing a common registration and ADS**.

We can illustrate the rationale for this recommendation by imagining a lengthy platoon composed of mixed-registration and mixed-ADS vehicles, and further that while in motion some, but not all, of the participating vehicles' automated driving systems determined they were no longer in an ODD. What would happen? The ensuing problem could certainly be solved, but a complex set of rules would be required to reach that solution; a set of rules understood not only by each vehicle's ADS, by each vehicle's fallback drivers (some of whom may be teleoperators), as well as by proximate non-involved vehicles, such as passenger vehicles that may be using an adjacent lane and may not have any level of ADAS or ADS engaged.

This solution engenders a subsequent problem, which is that this set of rules may require sufficiently complicated interactions, so that there will be diminishing returns as platoon size grows, and beyond a certain point new vehicles will elect not to join a particular platoon, meaning the value of a system to govern longer platoons will reach a ceiling. To date, our study of these matters, as discussed earlier in this chapter, points to the potential for extreme complexity and risk.

In the course of interviews taken while researching this work, we found that many of our interviewees were skeptical of platooning. According to one interviewee, "LCVs are better".[38] Certainly, LCV is a technology that exists now, which makes it attractive, while AV trucking is still on the horizon. All the same, we see merit in pursuing both. LCVs, though feasible, are fundamentally far more limited in their application. Because these vehicles must be physically connected, they can only travel to the same place for the same company (or purpose). Platooning does not share this constraint. Nonetheless we agree that platooning, beyond the achievement of short platoons with a single driver, deserves great care in implementation and full awareness of the risks that it poses.

On consideration of the complexity of the total system in which the use of platoons would be embedded, the risks of single points of failure, and the additional rules, technologies, tests, inspections, training, and due diligence needed, all while mixed with the foibles of human attention, we recommend that regulators view platooning beyond two-vehicle, common-registration (same fleet), same-ADS trains with great caution. The projected benefits of longer platoons will, in our view, be greatly eroded by their complications. In our view, the impact of train length on lateral string stability in situations of road debris, curved roads, poor surface conditions, etc. do not appear to be well understood.[39] Consider that a single excess crash could wipe out the net savings from many tens of thousands of tonne-kilometres of goods movement, even if savings are realized net of the cost of the additional system complexities. Regulators (and operators) should hold any commercial promises made for platoons beyond two vehicles to a very high evidentiary standard.

37 A heterogeneous mix of long platoons from multiple companies, jurisdictions, insurers, teleoperators, and ADAS/ADS software brands is difficult to risk-assess. Crash risk would be highest from increasing length, heterogeneous teleoperation and software, while settlement complexity would be exacerbated from having multiple companies with vehicles registered in multiple jurisdictions, and insured by multiple insurers. Crash risk might be mitigated by standards, rigorous testing, and frequent inspection, while settlement complexity could be reduced with harmonized registration and insurance rules. Studying and understanding these matters will be critical to permitting anything beyond the simplest, homogeneous two-truck configuration.

38 We note in passing that an argument we heard as to why LCVs are superior to platoons is that LCVs have greater fuel efficiency. While that may be the case now, we expect technological improvements to remove this advantage.

39 *String stability* as used here relates to the lateral stability of following vehicles on curves, during lane changes, and sudden steering inputs from the lead vehicle. This becomes harder to control with longer trains. Even if mishaps were rare, each one would be highly destructive. In other words, while benefits and risk both rise with train length, our understanding is that the configurations, lengths, and conditions at which the risks outweigh the benefits are neither well formulated nor understood.

To summarize, we assert that two-vehicle, uniform platoons operated by well-trained driver(s) and reliable companies, i.e., who will not cheat on inspections, weight and other restrictions, data delivery obligations, and additional aspects of system collaboration will be more likely to deliver the results that more than offset the difficulties of revising the regulatory environment to permit them. We further assert that the longer and more variable platoons become, the risks they offer will escalate faster and higher than their benefits.

10.2.5 A Note Regarding Determination of Consistent ODDS

In April 2022, at the annual plenary for ISO TC204 (technical committee for intelligent transportation systems), Working Group 14 proposed a Work Item dealing with the "ADS response to violations of ODD boundary conditions."[40]

This Proposed Work Item would generate a Technical Report (not a Standard) that sets out the principles for defining ODD boundaries, verifies the claimed ODD of an ADS, and determines factors for defining when to trigger a fallback driver based on an out-of-ODD condition, as described in Figure 5.

The Working Group's rationale for its proposal is that the Technical Report would provide guidance to engineers on the specificity of ODD attributes and alignment to system capabilities, to regulators on checks for ODD specifications and out-of-ODD condition, and to the establishment of WG14 ADS standards that involve the definition of ODD exit and boundaries.

At the time of this writing, ODDs are defined independently by each OEM for their respective ADS. This WG14 Work Item and its newly proposed Technical Report imply that compliant OEMs would follow a consistent definition for ODDs. If that were to be achieved, the maturation of ADS technology would permit closer operations of greater numbers of vehicles on a given infrastructure, or among platoons, dedicated to ADS operation. Without this development, mixing ADS brands within a single platoon would likely carry unknown and therefore, unacceptable risks.

40 One of the authors of this white paper attended this ISO Plenary meeting. The seven countries supporting WG14 are, in alphabetical order, Canada, Czechia, Germany, Korea, Sweden, UK, and USA

Figure 5 - Consistent Definition for ODD Boundaries

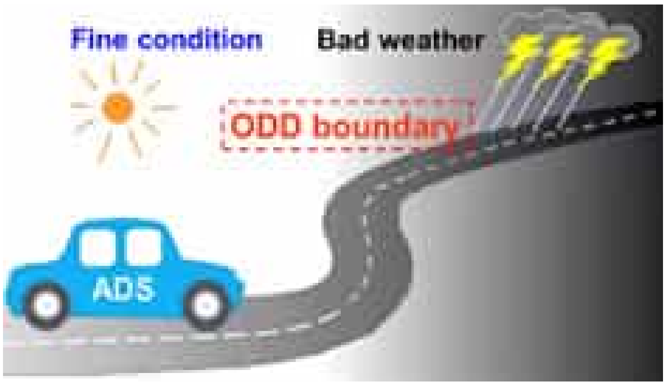

10.2.6 Levels of Human Engagement in Driving Automation

It is instructive to recall that the most recent revival of driving automation for civilian environments (passenger and goods) had its origin in the U.S. military's DARPA challenges of the mid-2000s. Out of this, the SAE3016 was developed with its definitions of driving automation capability, from which we, in this white paper, have isolated ADAS, ADS-H, ADS-V, and ADS-DV as the critical distinctions to understand the role, responsibility, and liability of the human driver.

The U.S. military has parallel terminology, which is easier to understand and is unburdened by the confusion with which popular journalism has encumbered driving automation and its jumbled focus on the SAE "Levels."

We include this comparison, here, in order to lead the reader to the observation that when the human is fully "out of the loop", as would be the case in ADS-V and ADS-DV engagements, that "machine errors and brittle automation can lead to cascading failures" [Cowen 2021]. It is such circumstances of failure that must be avoided, especially in the case of road trains (platoons) whether of passenger vehicles or trucks. This is not to say that having vehicles, or trains of vehicles without human drivers, is unworkable. Rather what it does say is that we must do everything to avoid "brittle automation". And this means we must be sure that there are backup safety systems and safe recovery methods for every aspect of these vehicles and platooning systems.

Table 5 - Comparison of SAE3016 and US Military Terminology for System Automation

SAE3016	DYNAMIC DRIVING TASK (DDT)	U.S. MILITARY		SYSTEM TASK
ADAS ADS-H	Human always present and immediately attentive	HITL	Human in the loop	Human has complete control to start/stop the automation
ADS-V	Human available when needed	HOTL	Human on the loop	Human has control only over autonomy planning
ADS-DV	Human not needed for DDT	HOOTL	Human out of the loop	Human has no control after launch. "…machine errors and brittle automation can lead to cascading failures"

11

Harmonization

11.1 At Issue

By harmonization, we mean the creation of common regulations among neighbouring jurisdictions, particularly as regards:

- Factors and measures of safety
- Rules of the road
- Vehicle user requirements
- Regimes for inspection and enforcement

While harmonization brings safety, operational and optimization benefits to all system users, it imposes political costs to regulating jurisdictions, both nationally and sub-nationally. Harmonization only works if all the parties participate in the standard that is set, which implies a particular responsibility to the national regulator to ensure its constituent sub-jurisdictions comply.

Harmonization is of special interest to the regulator for five reasons.

From a *global safety* perspective, it is best if all vehicles using a jurisdiction's roadways share a common understanding of driving rules and responsibilities.

From an *enforcement* perspective, it is best to have consistent enforcement guidelines and expectations from one jurisdiction to another, as regards rules of the road, the motive for and conduct of vehicle inspections, and vehicle operator responsibilities. Such consistency would have impacts on enforceability, driver training, licensing, and understanding, as well as court costs.

From a *training and licensing* perspective, it would advance road safety everywhere to ensure a common, consistent, and high level of user training and an assurance of the competence of licensed drivers.

From a *manufacturing* perspective—especially because of the complex requirements for sensors, software and software updates—it is best to have a single body of regulatory guidelines to incorporate into manufacturers' systems.[41] The presence of multiple rulesets among jurisdictions will not only make the production of these vehicles more complex and more expensive, but also risks the possibility of reduced safety if any software switches do not operate appropriately as vehicles cross jurisdictional boundaries. If each such crossing required a software refresh or a modified rules dictionary, automotive safety will necessarily be at risk; every unnecessary computation or telecommunication provides a new opportunity for error or failure.

From a *trade and commerce* perspective, it is best practice for a polity to harmonize its regulation of motor vehicles with its neighbours. Lack of harmonization will lead to friction, both for residents driving their vehicles abroad and for visitors bringing vehicles with them.

Notwithstanding all of these reasons, a polity also has a duty of care towards its residents and should be concerned that the regulations it promulgates are the safest possible. Consequently, a "lowest common denominator" approach is not the right way to harmonize, as this would make the least-safe jurisdiction the reference case. This constraint may limit the potential for complete harmonization.

41 This fact is in tension with the tendency in many polities for certification of a vehicle's suitability for use on the roads to be a national responsibility, while regulation of trained human drivers to be a subnational one. As automation increases, such polities will find the business of assuring compliance increasingly shifting from the latter order of government to the former.

11.2 Policy Analysis and Recommendations

11.2.1 Rules of the Road and "Over-the-Air" (OTA) Updates

Some observers whom we interviewed have proposed that harmonization is too difficult a problem to handle, and that this political issue may be solved via technology. On this view, the rules of the road, once set in each jurisdiction, could be programmed to be reloaded each time a vehicle crosses a jurisdictional boundary, via the vehicle's remote-Internet connection. Such a solution would allow each jurisdiction to set its own rules, to ignore any imperative toward harmonization, and to rely on the vehicles to abide by the regulations of each area through which they pass.

This technical solution is, in our view, far too risky to take seriously. It assumes no errors on the part of the human parties preparing updates; neither errors nor congestion in the telecommunication system; no meaningful lag as vehicles cross a boundary; no missing payments suspending software subscription updates; and no cyber-hacking. Put another way, if rules of the road are to be encoded into any ADS, that encoding and transmission must be six-nines (i.e., 99.9999%) secure, correct, and reliable. This is too high a bar to expect ADS to meet in real time at every jurisdictional boundary update.

Regardless of any agreement among jurisdictions regarding harmonizing rules of the road, it is still the case that these rules must be provided to vehicles according to a standard that has not yet been developed and in a way that is six-nines secure. We, therefore, recommend that the regulator **not legislate OTA updates as the *default* for cross-border changes to rules of the road**, though we agree with the universal understanding that we would rely on OTA updates for necessary adjustments to a vehicle's ADAS or ADS in the normal course of software maintenance and updates.[42]

The matter of specifying standards so that any jurisdiction can present its rules of the road to any vehicle within or entering that jurisdiction (again, regardless of harmonization) is a critical matter that cannot be handled independently by each province, state, or other subnational body. It needs to be an industrial or international standard. This is somewhat related to METR, as described in Chapter 5.2.

11.2.2 A Political Imperative

If a variegated technical solution is impractical and unsafe, a political solution is necessary. Regarding any driving automation and systems of management of driving automation, national harmonization is absolutely required. In large federal states, this may pose a challenge, as the national and sub-national governments often each have their own jurisdictions over driving regulation; nonetheless, the value of harmonization is such that each order of government should endorse a national standard, encourage other jurisdictions to agree, and make efforts to achieve this outcome. Additionally, we recommend **that the highest possible degree of harmonization be the paramount goal of driving-automation regulation for the next several decades, for all relevant parties within contiguous regions**: North America, Europe, Australia, Japan, East Asia, and so forth. We hasten to add that by this we do not mean that regulators dilute their commitment to safety, but instead that they pursue the harmonization of rules that take safe, consistent operation of ADS vehicles as the highest priority.

For the same reasons that Microsoft Word works identically in Peru or France, so too an ADS-enabled vehicle must operate identically in any two jurisdictions mutually accessible by road. As ADS embodies more and more regulatory information, harmonization of those regulations—including rules, user training, licensing, enforceability, responsibility, liability, insurability and data legibility—become critical.

There are two arguments against harmonization. The first is that it will take time and effort. We reject this argument on the grounds that the value of harmonization will more than offset these costs. The second is that not all jurisdictions may weigh all issues with similar levels of gravity. Unlike the first, we cannot dismiss this argument. Unless grappled with, harmonization could result in jurisdictions with laxer safety standards *de facto* imposing them on jurisdictions that would prefer more care be taken. Such a "race to the bottom" would be readily exploited by bad actors.

Preventing this latter outcome will require sustained effort. The most potent weapon that regulators may wield is the difficulty of arguing against propositions that will make roads and road users safe; both good conscience and the good opinion of constituents militate against accepting unsafe practice. To achieve high standards, regulators will have to persuade legislators that a given rule is the safest and therefore the best. Such persuasion will require good data, clear thinking, and a degree of eloquence. It is a task that demands committed energy and resources.

42 It may be ill-advised to rely on geographically-dependent OTA updates just as an ADS crosses a jurisdictional boundary. For example, if a person lived and worked on opposite sides of a jurisdictional boundary that was crossed daily, then either the data for both sides of that boundary would have to be retained and switched or otherwise reloaded several times a week, which would inevitably lead to a greater risk of errors. Such errors may come from telecommunication lags, non-payment of software licenses, telecommunication blind spots, software errors, data errors, cybersecurity issues, or inadequate testing. As much as possible, system updates should be managed while the ADS is otherwise idle to reduce the potential for a software error while a vehicle is in motion. A critical exception to this would be for the ADS to receive small and potentially continuous updates about things that pertain to the immediate ODD, especially as they may pertain to local weather, construction, or crash events. This is the approach taken by METR.

12

Data Collection

We recommend that regulators **determine which ADS-related operating data to collect, as well as the data's format, sampling frequency, protection, availability, permitted uses, and schedules for retention and destruction**.

Increasingly, it is possible for vehicles to record precise location and time-based events. Considering only dis/engagements of registered ADS-equipped vehicles, data has value to vehicle manufacturers, fleet operators, infrastructure designers, traffic planners, and government safety managers.

These data can be used to map infrastructural gaps in ODDs, improve and manage monetization schemes, and provide evidence of violations of operating contracts. A regulatory authority could track not only dis/engagements but also crashes, hard braking, or violations of the rules of the road. Presumably, any violation on the part of an ADS, and self-recognized by the ADS, would be one required by the circumstances of other traffic or of road infrastructure; such information would be critical for safety improvement programs.

While we stress the immense value these data will have for safety and order on roadways and for iterating ADS improvements, we recognize the hazards data collection provides. At an individual level, the movement of individuals, especially the origins and destinations of their trips, are and should remain private. At an aggregate level, datasets that are not widely shared offer an advantage to actors who possess them over those who do not, and datasets that are widely shared may, if not properly protected, allow bad actors to "re-identify" individuals.

Conversely, questions of whether and when particular ADAS or ADS features were active, are pertinent when settling questions regarding infractions, harm, and the assignation of liability. The public has an interest that bad actors face sanction for their misdeeds, but the public also has an interest in ensuring that the state does not have a presumptive right to know about the behaviour of private individuals. This is a perennial debate in liberal polities, and the arrival of always-connected ADAS and ADS adds a new complication.

Beyond these difficulties, it has become obvious to us in recent years that the public's relationship to data privacy, and simultaneously to the convenience that surrendering privacy can provide, is fraught. Casual data collection has the potential to excite strong emotions that could imperil the most well-intentioned of projects. For such an outcome to delay the arrival of ADS-equipped vehicles is a danger regulators should seek to avoid.

By way of illustration of the difficulties here, in Chapter 4, note that we did not recommend that the regulator insist upon the inclusion of an Event Data Recorder ("black box") to record the precise time and place of all attempted ADAS/ADS engagements and disengagements as a necessary feature of a registered vehicle. While a method to capture this data is vital, we cannot specify whether this data should be kept by the vehicle, within an external repository, or both. This matter is too consequential to make a simple recommendation. We recommend that regulators **consult with data and privacy experts in making their determination of which ADS-related operating data to collect, under what circumstances to make it available, and to which parties**.

Collecting data from ADS, if done carelessly, would pose a threat to OEMs, imperilling the entire driving automation project. Thus, we recommend that the regulator **determine a way to ensure sensitive collection and use of moving violation data**. It is important that OEMs believe that the collection of moving violation data will be used to help them improve their product, and not (or not only) to penalize them for such violations.[43]

For the regulator, privacy should encompass matters such as what data will be collected, its collection parameters, the policy governing its collection and use, how this data will be stored, analyzed, and destroyed, terms of access by third parties, and other relevant matters.

43 A related question is whether the fallback driver in an ADS-equipped vehicle is permitted to set the speed of the vehicle when the ADS is engaged. It seems to us a contradiction that the fallback driver, rather than the ADS, would be permitted to set the governing speed since the ADS is best equipped to understand its ODD. This is a conflict that is causing problems even today, where some drivers set their ADAS speeds above the posted speed limit, which is not only contrary to regulation, but also decreases overall safety, given that exceeding speed limits statistically increases the fatality rate in case of incident.

13

Loading, Unloading, Curb Space, and Parking

In our view, the matter of vehicle loading and unloading should be understood as integral to driving automation, and we recommend regulators **provide guidance to help cities understand and adjust to the required changes**. Parking and loading management will become considerably more complicated in the near-to mid-term while we have mixed automated and non-automated vehicle fleets, and absent such guidance, many municipalities will struggle unnecessarily and spend inappropriately.

Parking, whether to store a vehicle when not in use or to load/unload a vehicle while it is in use, is expected to change. Consider that if most vehicles are privately owned, then the parking problem remains a storage problem (as it is today). Alternatively, if most vehicles are service vehicles, such as robotaxis are projected to be, then the in-street parking problem becomes a loading problem similar to taxi queues at airports, except pervasive throughout cities.[44]

Regarding vehicle storage:

1. It is highly likely that a vehicle user, whether driver, fallback driver, or fleet operator, will be increasingly able to send or retrieve a vehicle to or from storage under direction of an ADS or teleoperator
2. Cities, employers, entertainment venues, etc. may rely increasingly on parking structures (garages or stackers) or offsite lots as parking facilities
3. Parking structures and offsite lots may become increasingly and eventually entirely automated, and as such may need new security and connected payment methods, rely more on stacking equipment, and rely more on "shuffling" algorithms to manage their ever-changing inventory, and some of these matters may require untested policy or regulatory attention
4. Parking structures that do not require human entry could be designed more efficiently, such as by featuring lower ceiling heights; these changes may have structural or fire implications that will require new forms of regulatory attention
5. Existing parking structures may be converted to other purposes, and other-purposed structures may be converted for parking use; while these conversions may require regulatory attention, it is likely that building codes as they stand can address the matter with simple review

ISO 5206-1 addresses some of these issues, notably excluding structural and fire matters. The standard is in its final draft for publication in 2023.

For loading and unloading:

1. It is certain that ADS-equipped passenger transport and goods delivery at scale will require dedicated pick-up and drop-off zones (PUDO; loading zones) in many locations
2. It is highly likely that ADS-equipped vehicles will require an orchestration system to manage booking or reservation for loading and unloading, especially in any area where there will be competition for access; this would be a "ground-traffic control system"

In our view, the problem of parking-for-loading/unloading passengers and goods will be a larger and earlier problem

44 Some have suggested that driving-automation technology, combined with cloud computing and the mobile internet to permit sharing and booking of private assets, could permit a privately-owned vehicle to generate income as a robotaxi when not being used by the owner. If this suggestion proves true, there may exist a hybrid owner/user option between privately owned/used and fleet owned/operated; for example, a private car club in which several families co-owned a smaller number of vehicles (perhaps 100 families and 20 cars).

than the parking-for-storage problem. Again, we note that ISO has standards that do, or will, address these matters. ISO 4448-4 addresses loading and unloading orchestration issues, using ISO 5206-1 as its spatial basis.

Because of this new complexity of parking requirements and circumstances, there may be a considerable learning curve for each city to work out appropriate solutions. Any shortfall in these local solutions will degrade the benefits of ADS technology in those locations. Hence, a body of guidance such as proposed in ISO DTS 4448-4, short of regulation, would be helpful.

14

Public Mobile Robots

Public mobile robots differ significantly from automated passenger and road vehicles due to their unique characteristics, such as smaller size, lower speeds, greater diversity of operating locations and purposes, and close interactions with non-involved humans. Hence, it is crucial to gain a deeper understanding of them before addressing the need for regulatory measures.

Since 2005, significant funding, innovation, and media attention have been directed toward developing automated vehicles, including personal household vehicles, robotaxis, passenger shuttles, goods vans, and heavy-duty trucks.

In contrast, the public mobile robot, such as the last-mile delivery robot, has received comparatively little attention and investment—three orders of magnitude less. This disparity is expected to change as multiple factors drive demand for these devices.

1. The tendency for consumption and convenience has accelerated e-commerce over the past decade. Further boosted by Covid-19, last-mile, short-notice delivery demands, and a concomitant labour shortage have increased demand for robots, initially in factories and warehouses, and now beginning to reach the neighbourhood sidewalk
2. The growth of e-commerce and the resulting increase in last-mile delivery needs, combined with labour shortages and congestion, have led to an increase in interest in smaller vehicles like e-bikes and robots
3. Technological advancements, including in software, AI, IoT, geographic data and algorithms, as well as improvements in mechatronics and warehousing, have made robots more efficient, cost-effective, and desirable. With last-mile delivery being a major supply chain expense, automating it is becoming an attractive opportunity for robots
4. Governments, such as municipalities, face costs related to the cleaning, maintenance, and surveillance of public spaces. These costs can be reduced by deploying robots for these services. The aging population and labour shortages are increasing the demand for public mobile robots in several cities
5. De-globalization and re-shoring also drive the need for robotic innovation, making deploying robots in public spaces more feasible by association. The vision of human-scale robots effectively navigating public spaces is rapidly becoming a reality

14.1 At Issue

Accessibility: Universal access rules must be reviewed to ensure PMR technology does not negatively impact existing accessibility measures. To ensure that pedestrians with impaired sight or hearing can understand the intentions of nearby robots, it will be critical for PMRs to use standardized sound, light, gestural, and/or haptic communication signals.

Complexity: The deployment of these robots raises unique social, traffic, safety, usage, and governance issues that differ from those of automated cars and trucks.

Crosswalks: Road crossings also require proper jurisdictional enforcement. This becomes even more complex when bike lanes are involved. Should appropriate guidelines be national, regional, or municipal?

Engagement: Stakeholder engagement from various interests, including municipalities, planners, accessibility advocates, merchants, logistics providers, and roboticists, is crucial in the policy-making process.

Safety: Using robots on sidewalks and bike lanes raises both safety and regulatory challenges. While these areas are designed for slower speeds, they are often used by vulnerable road users such as pedestrians and cyclists, making the presence of robots a new form of potential hazard.

Standards: ISO draft technical standards are underway. It is important to carefully consider the impact of this technology on safety and accessibility and engage the relevant stakeholders to make informed policy decisions well before this technology reaches scale.

Uniform Bylaws: Municipal regulators should seek and consider national or international standards to ensure consistent and safe rules for using these robots.

Urban Planning: Inadequate pedestrian infrastructure and the potential for robots to exacerbate existing problems further highlight the need for proper regulation.

14.2 Literature Review

Literature regarding public mobile robots from a social or deployment perspective is nascent, rarely touching on safety and operating standards.[45] A majority of this extant material focuses on last-mile, small-package, delivery robots, or security and surveillance robots. It largely ignores other applications such as sweeping, removing litter, ploughing snow, de-icing, or metering parking violations. To date, there is very little research beyond the development and testing of such devices, and some of its reporting is influenced by marketing bias.[46] While there is a natural assumption that it is good for there to be machines that improve a city's maintenance capability, few studies show the social, safety, sanitary, or cost benefits of such devices because the technology remains underdeveloped and under-deployed. Surveillance robots operating in municipal spaces may be an exception to this.

The literature regarding delivery robots, however, is more extensive and worth consideration.

The human-factors engineer Michael Clamann has written on the interaction of pedestrians and cyclists with delivery robots. He and co-author Meg Bryson write about personal delivery devices (PDDs, a specific class of public mobile robots):

> While developers are testing more deployments in more states in the wake of new legislation, several challenges exist, and the real impacts of PDDs will remain unknown until they are deployed in large numbers. Early deployments appear to operate in settings with high-quality [well-maintained] physical infrastructure and limited operational ranges, like university campuses. While many states require lights on PDDs, many locations opt to only operate during daylight hours. PDD mobility remains well behind human mobility in terms of being able to negotiate curbs, thresholds, stairs, and damaged and cluttered walkways, as well as the ability to safely interact with other users of the walkway space. Like automated vehicles, broad deployment of PDDs depends on regulatory decisions, public trust and acceptance, and technology readiness.[47]

Alanna Coombes, UCLA, was involved in early trials for delivery robots in London (UK), circa 2018. Concerned with matters of rights to access, use, and enjoyment of public space in cities, she reaches beyond matters of safety and commercial optimization into the full purpose and expression of footways and sidewalks in our cities and how service robots might influence these:

> To thrive we need community, business and political agreement on who has rights to the kerb and footways. In turn, these rights need to be turned into clearly defined priorities that meet the needs of citizens, including those traditionally excluded, and businesses. Public space, including the kerb and footways should be designed for community and artistic expression and livability. These vital public spaces – like the city centers in which they exist – need to adapt to the needs of current and future generations, addressing their economic, social, community needs and their wellbeing.[48]

One of the most important early writers regarding public mobile robots is Kristen Thomasen, a leading Canadian expert in robotics law and policy with the Peter Allard School of Law, UBC. She writes:

> Public space is a complicated socio-legal concept. … legal definition is one of several factors that combine to render a space public, or not. Whether members of the public can access this space, and how conduct within that space is regulated, also contributes to its public nature. This paper has considered the impact that robots and robot regulation will have on the public nature of public spaces. [It] has sought to emphasize that lawmakers need to be careful and explicit about how they regulate robotic systems that operate in public spaces, because by regulating robots, lawmakers may also be implementing a particular vision of public space that renders that space more or less public to different individuals and communities. … where a robotic system serves to make a space more accessible, lawmakers should be cautious to avoid providing differential access to that space through the regulation of that robotic system… lawmakers should resist any arguments by users or manufacturers of robotic systems that public space, by virtue of its public nature, should be freely available for the use of robotics. Such an approach threatens to privatize and commercialize public spaces in ways that would exclude people, and would entirely overlook the already exclusionary impact of the colonial laws and systems operating in these spaces.[49]

45 The Urban Robotics Foundation, of which one of the authors of this white paper is a cofounder, is a Canadian non-profit established in 2021. The URF, as the project leader for the ISO standard "Public mobile robots," is poised to redress this gap. These international standards are designed to coordinate robotics policy across municipalities.

46 One exception is Loke, S., and Rakotonirainy, A., (2021)

47 Clamann, M., Bryson, M. (2021).

48 Coombes, A., Grush, B. (2022).

49 Thomasen, K. (2020).

Lastly, Cindy Grimm and Kristen Thomasen provide advice regarding "at least five stakeholders that need to come to the table when discussing public mobile robot regulation."[50] These are:

1. Robot technicians/technical experts to discuss what is feasible from a technology standpoint
2. Robot company business financial/marketing, to discuss business models, cost-effectiveness, marketing (to date, these stakeholders have dominated discussions with U.S. regulators)
3. City planners, to discuss city infrastructure, development plans, and considerations of zoning and congestion
4. Residents, to discuss public space uses and privacy concerns
5. City businesses, to discuss the potential use of this technology

14.3 Jurisdictional Scan

14.3.1 Estonia

Estonia is the birthplace of the delivery-robot firm Starship and arguably the birthplace of the public-space delivery robot. Starship's robots operate in at least two Estonian cities on sidewalks. Their largest Starship fleet, reportedly comprising 200 devices in 2021, operates in Milton Keynes (UK), a city whose urban form happens to be highly suitable for these devices. In a presentation at the UN on March 8, 2022, a company spokesperson boasted that since its inception, it has conducted 3 million deliveries with an average trip length of 2km. It also claimed that its devices cross 140,000 streets each day, or (on average) three per second.[51] Street-crossing is the most critical safety element for all such devices and something that any transportation authority should consider closely. Note that the bulk of this achievement has been on college campuses that provide fewer challenges with regard to automotive traffic and intersections and greater acceptance from a generally younger population sharing the footways.

14.3.2 Japan

As of April 2023, Japan will release a revised Road Traffic Law that adds "remotely operated, small vehicles" to the list of definitions of "pedestrian." This update details numerous rules and obligations for using footways and other permitted infrastructures. While similar to legislation from various U.S. states that analogizes these devices to pedestrians, this revised law cites several "rules of the road" and explicitly recognizes the urgency for an enforcement officer to be able to stop or move a remotely operated vehicle to prevent danger and to eliminate interference with traffic.

While this is the most recent national ruling as of this writing, it is likely that many hundreds of traffic jurisdictions will incorporate such revisions in their traffic laws over the next decade. As these legislative changes take hold, they will become more insightful and likely far more rule-bound with regard to the governance that will be required. This will especially be the case in those jurisdictions that admit a significant number of these devices from multiple operators, providing multiple services around the clock.

14.3.3 South Korea

By contrast, within **South Korea's** traffic laws (as of March 2022), autonomous robots are classified as 'vehicles,' which prohibits their access to sidewalks, crosswalks, and parks. In January 2022, South Korea's Office for Government Policy Coordination released a statement indicating its commitment to revise these laws by 2023, two years ahead of its original schedule. In the meantime, "outdoor autonomous delivery robot" programs operate within a regulatory sandbox. It is speculated that the updated Intelligent Robots Development and Distribution Promotion Act will establish a robot management system, including a robot safety certification system, and further define liability laws. Regulations for outdoor autonomous robots are expected to be relaxed in three phases.

1. Initially, delivery robots would be allowed to use elevators for greater indoor services, expected typically within hospitals and hotels
2. Then, to travel on pedestrian sidewalks and crossings
3. Lastly, to use bicycle paths for travel at bicycle speeds

The latter change will require a revision of South Korea's road traffic laws. The sources we accessed that discuss these changes are in the English-language business press; in this context, they do not discuss the consequences to pedestrians, cyclists, or pedestrian and cycling infrastructure.

14.3.4 United States of America

20 U.S. states have passed legislation permitting "personal delivery devices", a specific type of public mobile robot, to transport goods over short distances—usually two or three kilometres.[52] This legislation merely sets a general frame for their size, speed, lights, brakes, insurance and guidance, such as yielding to pedestrians and similar behavioural issues, leaving more detailed issues to the municipality.

In **Michigan**, public mobile robots are defined as devices that are no more than 40 inches wide, weigh less than 400 pounds (including cargo), have a maximum attainable speed of 25 miles per hour or less (roadway) or on a sidewalk with a maximum speed of 10 miles per hour, is equipped with an automated driving system and is used to transport goods or perform services on a sidewalk or other areas open to

50 Grimm, C.M., and Thomasen, K. (2021).

51 This information was presented to a UN panel in March 2022. It is consistent with many other sources of online information we found available over the past two years about the subject firm, so we judge it to be reliable. What was not reported by this company are any issues they have had with pedestrians. There have been at least two incidents independently reported in the press, but they appear to be minor (no injuries; neither were crashes). Consider also the number of roadway crossings indicated given only a few hundred robots. In the future, the number of robots, and therefore the implied number of roadway crossings, would be several orders of magnitude higher. This leads to questions of road-crossing safety.

52 At the time of writing, the website Personal Delivery Devices (PDDs) Legislative Tracker provides an excellent overview of state-by-state legislation in this regard.

pedestrian traffic or a highway or street, but not to transport a human operator or passenger.

Pennsylvania defines public mobile robots as devices that are manufactured for transporting cargo or goods, are operated by an automated driving system or a driving system that allows remote or autonomous operation, or both, and weigh 550 pounds or less without cargo or goods. The maximum permissible speed is 25 miles per hour.

Virginia's SB1207 definition of public mobile robots is typical of many U.S. state legislative definitions. It reads as follows: "Electric personal delivery device" means an electrically powered device that (i) is operated on sidewalks, shared-use paths, and crosswalks and intended primarily to transport property; (ii) weighs less than 50 pounds, excluding cargo; (iii) has a maximum speed of 10 miles per hour; and (iv) is equipped with technology to allow for operation of the device with or without the active control or monitoring of a natural person."

U.S. states set various weight limits, with **Virginia** being the most constrained at 50 lbs and **Michigan** and **Pennsylvania** being significantly more generous, with a maximum gross weight of 400 and 500 lbs, respectively.

These states have restricted the transport of humans on these devices. Some permit PDDs to travel at a maximum speed of 25 miles per hour, and the devices have to yield to pedestrians and cyclists and are only permitted to overtake in a manner that is safe for the human actor. All states require the PDD to be monitored by a human but permit the operation of the device by ADS or a human, so long as that person is over the age of 16 years. Furthermore, the operator is required to be able to take immediate control of the PDD in the event of an ADS failure.

In all states, there is a requirement for the PDD to display the manufacturer's details and also possess suitable brake and light fittings.

Michigan and **Virginia** do not have specific training requirements as yet; however, licensing is through the state transport department and is valid for three years. **Pennsylvania** requires PDD operators to be trained in accordance with the manufacturer's regulations.

All states require liability and insurance coverage. **Michigan** requires the operator of the PDD to maintain liability coverage on each device of not less than $250,000 for damages caused by the operation of the device. Furthermore, the state stipulates that the operator of an automated delivery device is liable in the event of personal injury or property damage caused by the operation of the device. **Pennsylvania** and **Virginia** require the authorized entity to maintain an insurance policy that includes general liability coverage of not less than $100,000 per incident for damages arising from the operation of the PDD, and further that the state and municipalities are immune from suit from personal damages caused by these devices.

All these states identify the PDD as 'pedestrians' and, as such, are subject to all laws and regulations governing pedestrians.

14.4 Regulation

The regulation of public mobile robots is a topic of ongoing discussion, with safety being a primary concern. The reception of PMRs has varied, with some cities embracing them while others have banned their use. PMRs, such as those used for food or grocery delivery, have been banned in San Francisco, Toronto, Ottawa, and Tel Aviv, while they have been well received on many university campuses and in Tallinn, Milton Keynes and at least four other UK cities and towns. In cities such as Detroit, Miami-Dade, Pittsburgh, and San Diego, the adoption of PMRs has been lukewarm—sometimes more the result of inadequate pedestrian infrastructure as of the service value they offer.

There are several potential sources for hostility: unfamiliarity with the technology, fear of job loss, or the overuse of the trope of the "killer robot" in popular media. Another is municipal governments' troubled experience with ride-hail and e-scooter companies, which had begun operation in many cities without consulting regulators and had antagonized residents. To avoid these outcomes, some cities have created obstacles to PMR trials, sometimes out of proportion to any reasonable cost-benefit analysis, in addition to the few that have flatly banned robot operations in public spaces. At least one permitted a food delivery pilot, but in a location lacking customers or retailers. With neither shippers nor receivers, there would be nothing commercial for delivery robots to do.

PMRs used for surveillance purposes have generally received a better response as they operate in areas where security is a priority. PMRs for police security have had a mixed reception. One such machine that could assist the police with reconnaissance and bomb disposal was trialed in San Francisco in late 2022. This deployment was curtailed within days under protests concerning its designed potential to carry explosives. Robotic operation via remote control is permitted but only if not weaponized.

The example of using robots in police work is one of the innumerable instances where the potential for unintended consequences weighs heavily against the intended value. The value of making police work safer is almost always seen in a positive light, but the potential for mixing weapons with robots must be carefully considered.

Many early PMR trials that were permitted were put in place without proper regulation, highlighting the need for thoughtful oversight. What would be a starting place to pre-regulate sufficiently to be able to trial this technology while waiting for detailed regulations to be prepared? We suggest:

Regulating the combination of speed and mass. Robots pose a danger to nearby people if they are travelling fast enough and heavy enough. In the event of a collision with a human, an injury would be possible. Regulating the level of this kinetic energy by specifying both weight and speed to fall below a given threshold would reduce risk and uneasiness, especially if the threshold was calibrated to a particularly vulnerable sidewalk user, such as an elderly person or a child. For this reason, robots should be discouraged, at least initially, from operating in bike lanes because the speeds they could reach there would make them more dangerous for the vulnerable road users that they would encounter, whether cyclists or jaywalking pedestrians.

Requiring a fail-safe. At present, responsible roboticists and operators include a variety of fail-safe mechanisms in their PMRs. They do this as a common-sense business strategy and to be good corporate citizens, as there are often no specific regulations for this. Absent this, how can a municipality be certain of the safe operation of the devices to be permitted? How can an insurer calculate premiums?

The harm of a serious failure would spread to all stakeholders—not just the bystander and PMR operator, but to the shipper and the PMR industry as a whole. Because the industry is nascent, errors which cause harm, or worse, to a human would have considerable repercussions, just as an automated test vehicle struck and killed a cyclist in Tempe, Arizona, in 2018. A fatality involving a PMR, however unlikely, is something that every stakeholder must avoid. For this reason, all responsible operators seek thoughtful regulation to prevent making everyone worse off.

Regional, sub-national, or national regulation is most helpful. If every city requires a pilot or trial, the ensuing costs and duplication of effort will delay or prevent the value achievable from this technology. Regions and neighbouring cities should find ways to collaborate to establish safe operations.

14.5 Discussion

In their brief history, automated mobile robots (AMRs) have been used successfully in factory and warehouse operations—first to move goods, then to stack and pack. In parallel, AMRs spread to farms and mines. In these repetitive settings, AMR's tirelessness and precision have made them ideal substitutes for human labour. At least, they have done so in a controlled, monitored space, as befits expensive machines with the potential for considerable harm if used or handled carelessly.

Since approximately 2015, AMRs have matured into specialized PMRs to break out of industrial and warehouse environments into public, pedestrianized spaces. Today, PMRs operate in hospitals, restaurants, university campuses, and airports as tools for delivery, maintenance, and surveillance. There is already a fledgling industry of PMRs entering less-structured municipal footways and bikeways to deliver packages and food, sweep streets, plough snow, pick up litter, and even write parking tickets.

Sized for footways and bikeways and initially radio operated as large toy cars, mobile robots are being equipped with cameras, LIDAR, communication tools, and intelligent software developed for automated passenger vehicles. Some are claiming SAE Level 3 and 4 automation capabilities. This innovation is in its earliest stages of maturity and enjoys far less government oversight for the confirmation of automated capabilities compared to that dedicated to driverless passenger vehicles.

Their variation, versatility, and capabilities are expanding rapidly. They are evolving in parallel with IoT technologies. They are being enabled to walk, climb stairs, open doors, place packages into street lockers, and communicate their intentions to humans. Still, the application of mobile robots in public space is only in its infancy.

Within our cities, the valuable role these small electric devices can play is evident. A 35-kg electric device making a 5-kg delivery is a vast improvement over a 1,600-kg car making that same delivery, even if both device and car are powered by sustainably-generated electricity (if the car is not electric, the comparison is even more favourable). A PMR's ability to quietly remove litter while people sleep, provide safe area surveillance, spread carefully measured salt on walkways, or clear blocked storm drains, could vastly improve the quality and safety of public space for minimal cost and disruption. The far-reaching potential of this technology is evident.

14.6 Sharing This Space with Existing Users—Who Has Which Rights?

Taking advantage of this potential is not straightforward. Early AMR applications were deployed within controlled workspaces where nearby humans were trained to collaborate with them. PMRs are powered devices sharing space with humans that are not trained collaborators. These spaces have previously been reserved for pedestrians who are not familiar with, and may not always be prepared to interact with, PMRs. Complicating matters, these spaces also support other participants, such as pets, baby carriages, and wheelchairs. Additionally, some PMRs are intended for bike lanes that are designed for through-traffic from cyclists—themselves vulnerable road users. And some of these footways and bikeways may be narrow, cluttered, poorly designed, or poorly maintained. Footways especially contain fire hydrants, trees, newspaper boxes, garbage bins, and retailers' signs and wares. Such spaces are organized for strolling, window shopping, waiting for the bus, watching street performances and now for restaurant dining. Some are used for sitting, begging, or sleeping.

Mobile service robots can clearly bring enormous advantages while also testing previously settled matters regarding shared rights-to-use. What rights should retailers and restaurant operators have to use robots to deliver goods or groceries on these footways and bikeways? What about senior residents who wish to receive deliveries or demand improved walkway maintenance? Or pedestrians who have mobility, sight, or hearing losses? What about footways that are already inadequate for existing traffic in their dimensions or conditions? How should these devices interact with pedestrians at intersections, already the location of a majority of pedestrian and cyclist fatalities?

14.7 Standards Are Under Development

In 2019, the International Organization for Standards (ISO) approved a project to draft a new standard for managing the real-time queueing and management of loading and unloading of robotic vehicles, such as automated taxis for passengers and automated goods vans at the curbside.[53]

53 In the interests of full disclosure, we note that one of the authors of this white paper leads the drafting project for this technical standard: ISO DTS TC204 WG19 4448 Public mobile robots.

In 2020, this project was expanded to include robots on the footway or sidewalk—the domain of the pedestrian. This meant consideration of robot behaviour: how robots should give way to pedestrians, how they should communicate their intentions to blind or deaf pedestrians, how they should use crosswalks, etc.

By 2021, the scope of the draft standard had further expanded. How should robots enter and leave bike lanes? How should they behave while passing busy bus stops? From a traffic management perspective, how could their numbers be limited within each block face or during peak hours? As of 2022, elements for robot safety and various aspects of readiness certification have been added.

Even considering all this, there is still a deeply profound issue for these robots in our cities. What is being introduced are small, motorized machines that can roll, walk and flow in and out of our footways, bikeways, and road shoulders—traversing any infrastructure they are equipped and permitted to use. Likely, over the next decade, these devices will become far more coordinated and spatially capable than active transportation devices are today, as well as remaining cleaner, smaller, and quieter than our current motor vehicles. They will almost certainly improve dramatically in capability and environmental suitability and–we anticipate–become more spatially nimble than most pedestrians.

We want all of the work and environmental advantages this technology offers, but as with everything humans invent, there are unintended consequences. Consider a time in the current decade in which:

1. a large variety of robots,
2. that have multiple purposes,
3. which are each independently operated,
4. by multiple independent operators,
5. performing maintenance, delivery, and monitoring activities,
6. each on independent and asynchronous schedules,
7. all competing within a common public space with each other and with human pedestrians.

We could face a traffic management problem far greater and more complex than our current urban traffic management problems.

An easy solution is to ban these devices. Some cities have already taken this step–at least temporarily—as discussed earlier in this chapter. We cannot recommend outright banning. There are too many advantages these devices bring, but there are also too many risks to leave them ungoverned. And banning seems not to work as a permanent solution. San Francisco banned delivery robots in 2017 but has since cautiously lifted that decision to admit surveillance and security robots.[54] More interestingly, at least one delivery operator is now operating in nearby Pleasanton in the Bay Area. Similarly, three or four different ones operate in small towns in the Los Angeles area, such as Santa Monica and West Hollywood.

54 https://www.bloomberg.com/news/articles/2017-05-19/the-sidewalk-robot-resistance-begins-in-san-francisco

The deployment trend is first to consider the easy case of college campuses, then the somewhat more difficult case of retail deliveries within a small downtown. Santa Monica is only about 90,000 people; West Hollywood is 35,000. Regardless of location or population size, neither cities nor the robotics and last-mile logistics industries can operate without agreed standards. Such standards must incorporate matters for safety, data, governance, machine behavior, and traffic orchestration. There will be ample reason for monetization standards, as well.

Ideally, international standards should inform national model codes. Such national codes would inform jurisdictional legislation, which in turn would inform local municipal bylaws. To date, this approach has not prevailed. In approximately 20 American states, a piecemeal mix of legislation has been passed—in our view, much of it hasty and inadequate. This has prodded several U.S. municipalities to instigate pilots and trials, which will almost certainly result in local bylaws. This will make matters more complex for each municipality and each logistics, food delivery, and maintenance operator.

14.8 Model Codes

Figure 6 below illustrates an ideal cascade from standards through model codes and legislation to local bylaws. For a common set of standards to percolate through in this fashion, some time would be needed. Work started on the ISO 4448 standard in 2020, and it is currently in draft form. Estonia (*Traffic Act*), Japan (*Revised Traffic Act*) and South Korea (*Road Traffic Act and Intelligent Robot Act*) have each released revised traffic acts. There are several state legislature bills passed in the United States, as mentioned above. There are several cities that have municipal bylaws on their books. The two in Ontario, Canada both temporarily ban the devices while they are being studied.

Figure 6 - From Standards to Bylaws- the Ideal Progress for Governance of Public Mobile Robots

Over the next few years, robot operators will contend with a growing variety of operating processes in many jurisdictions. To address the likely gap in harmonization, we recommend that regulators **consult with accessibility, logistics, municipal, planning, and robotics experts to determine a reasonable approach for its cities in any case where regional or national guidance is unavailable**. To this, regulators may wish to add residents and city businesses, as suggested by Grimm and Thomasen in a paragraph above. Together, such consultation would provide guidelines and certification methods to ensure that robotic passenger and goods systems are safe, managed, and result in improved livability for a community.

In a perfect world, guidance and legislation would proceed in this order, with each order of government taking its cues from the one preceding. In practice, however, the smaller the government, the fewer resources it has available to devote to such questions, while the larger the government, the harder it is to proceed proactively. Consequently, it is often subnational governments that will lead in these matters, bringing regional and sometimes national governments along with them. Local governments that seek to avoid preemption may need to develop regulatory approaches proactively.

14.9 ISO 4448—Draft Technical Standard for Public Mobile Robots

Working Group 19 of the ISO Technical Committee for Intelligent Transportation Systems (TC204/WG19) has commenced drafting standard 4448, "Public mobile robots and automated pathway devices," for

1. Vehicles engaged in the pick-up and drop-off of passengers and goods at the curbside
2. Automated and robotic mobile service and logistics vehicles
3. Independent service robots using walkways, bikeways and roadways

This draft technical standard is of value to any jurisdiction considering the use of pedestrian or cycling infrastructure by mobile robotic devices. Among its multiple parts, ISO/4448 addresses behaviour, orchestration, municipal readiness, safety, enforcement, and monetization.

Among these core issues are intersection and crosswalk safety, congestion management, consistent distancing, robot-to-human intention signalling, emergency responses and many more topics. Consider that in a typical town or city, a two-km robot delivery (4 km round trip) can expect to cross approximately 22 streets. That implies critical safety and traffic control considerations for any city where these devices begin to scale.

The ISO 4448 series, when completed, will comprise a set of terminology, guidelines, and real-time procedures for the safe coordination of operations at the curbside, on pathways, and the integrated use of both curbside and pathway. The data and communications standards being defined in this series of deliverables are intended to enable carefully defined (mapped) and expanding areas of cities to manage an arbitrary number of vehicles and vehicle varieties operated by any number of operators (public, commercial, and private) for their various service and economic activities.

We recommend that cities start preparing their rule book now.

15

Bibiliography

Aramrattana, M., et al (2021). "Safety and experience of other drivers while interacting with automated vehicle platoons", Transportation Research Interdisciplinary Perspectives. Retrieved from https://www.sciencedirect.com/science/article/pii/S2590198221000889?via%3Dihub. Accessed 14 March 2023.

Autonomous Vehicle Industry Association (AVIA) (2021), "Explainer: Driver-Assist Tech is Not an Autonomous Vehicle". Retrieved from https://theavindustry.org/resources/testimony/explainer-driver-assist-tech-is-not-an-autonomous-vehicle, 3 October 2021.

Calo, R., Froomkin, A.M., Kerr, I., (2016). "Lawyers and Engineers Should Speak the Same Robot Language", in Robot Law, ed. Ryan Calo, A. Michael Froomkin, and Ian Kerr. Elgar Publishing, 2016.

Chatham House. "Chatham House Rule". Retrieved from https://www.chathamhouse.org/about-us/chatham-house-rule. Accessed 14 March 2023.

Chrisrussolaw.com. (2014). Florida's Autonomous Vehicle Statute - Russo Law | Personal Injury, Car Accidents, & Medical Malpractice. [online] Retrieved from https://chrisrussolaw.com/blog/car-accidents/floridas-autonomous-vehicle-statute/. Accessed 27 March 2022.

Clamann, M., Bryson, M. (2021) Sharing Spaces with Robots: The Basics of Personal Delivery Devices. An info-brief from the Pedestrian and Bicycle Information Center. Retrieved from https://www.pedbikeinfo.org/cms/downloads/PBIC_InfoBrief_SharingSpaceswithRobots.pdf. Accessed 12 February 2023.

Clamann, Michael (2021). "Sharing the sidewalk with robots". This is a video of an in-person talk, and contains a link to a "personal delivery device tracker" of US Senate Bills regarding sidewalk robots. Retrieved from https://www.enotrans.org/event/webinar-sharing-the-sidewalk-with-robots/. Accessed 12 February 2023.

Cornell Law School Legal Information Institute (May 2020). "Reasonable Reliance." Retrieved from https://www.law.cornell.edu/wex/reasonable_reliance. Accessed 14 March 2023.

Coombes, A., Grush, B. (2022). Digitization, Automation, Operation, and Monetization: The Changing Management of Sidewalk and Kerb 2000-2025. In J. R. Vacca (Ed.). Smart Cities Policies and Financing: Approaches and Solutions. Elsevier.

Cowen, M., Williams, R., Cherry, D., (2021) Is Human-On-the-Loop the Best Answer for Rapid Relevant Responses (R3)? Retrieved from https://www.japcc.org/essays/is-human-on-the-loop-the-best-answer-for-rapid-relevant-responses/. Accessed 15 February 2023.

Florida Legislature. (2021, January 1). Statutes & Constitution: View Statutes: Online Sunshine. Retrieved March 27, 2022, from http://www.leg.state.fl.us/statutes/index.cfm?App_mode=Display_Statute&Search_String=&URL=0600-0699/0627/Sections/0627.749.html. Accessed 14 March 2023.

Grimm, C.M., and Thomasen, K. (2021) On the Practicalities of Robots in Public Spaces. Retrieved from https://werobot2021.com/wp-content/uploads/2021/08/GrimmThomasen_Sidewalk-Robots.pdf. Accessed 15 February 2023.

Holland-Letz,D., Kässer, M., Kloss, B., and Müller, T. (2021). Mobility's Future: an Investment Reality Check, McKinsey & Company.

ISO-METR. (2021). Management of Electronic Transport Regulations. [online] Available at: https://iso-tc204.github.io/iso24315p1/index.html. Accessed 5 January 2022.

KPMG, (2020), Autonomous Vehicle Readiness Index (AVRI). Retrieved from https://home.kpmg/xx/en/home/insights/2020/06/autonomous-vehicles-readiness-index.html. Accessed 14 March 2023.

Li, S., Sui, P.-C., Xiao, J., & Chahine, R. (2019). Policy formulation for highly automated vehicles: Emerging importance, research frontiers and insights. Transportation Research Part A: Policy and Practice, 124, 573–586. https://doi.org/10.1016/j.tra.2018.05.010. Accessed 14 March 2023.

Loke, S., and Rakotonirainy, A., (2021) The Automated City: Internet of Things and Ubiquitous Artificial Intelligence.

Manser, M., et al (2019) Driver Training Research and Guidelines for Automated Vehicle Technology. https://vtechworks.lib.vt.edu/handle/10919/95178

Mordue, G., Yeung, A., & Wu, F. (2020). "The Looming Challenges of Regulating High Level Autonomous Vehicles." Transportation Research Part A: Policy and Practice, 132, 174–187. Retrieved from https://doi.org/10.1016/j.tra.2019.11.007. Accessed 14 March 2023.

National Cooperative Highway Research Program (NCHRP). (2018). Implications of connected and Automated Driving Systems, Vol. 1: Legal Landscape. https://doi.org/10.17226/25296. Accessed 14 March 2023.

National Cooperative Highway Research Program (NCHRP). (2018). Implications of connected and Automated Driving Systems, Vol. 2: State Legal and Regulatory Audit. https://doi.org/10.17226/25294. Accessed 14 March 2023.

National Cooperative Highway Research Program (NCHRP). (2018). Implications of connected and Automated Driving Systems, Vol. 3: Legal Modification Prioritization and Harmonization Analysis. https://doi.org/10.17226/25293. Accessed 14 March 2023.

National Cooperative Highway Research Program (NCHRP). (2018). Implications of connected and Automated Driving Systems, Vol. 4: Autonomous Vehicle Action Plan. https://doi.org/10.17226/25292. Accessed 14 March 2023.

National Cooperative Highway Research Program (NCHRP). (2018). Implications of connected and Automated Driving Systems, Vol. 5: Developing the Autonomous Vehicle Action Plan. https://doi.org/10.17226/25291. Accessed 14 March 2023.

National Cooperative Highway Research Program (NCHRP). (2018). Implications of connected and Automated Driving Systems, Vol. 6: Implementation Plan. https://doi.org/10.17226/25295. Accessed 14 March 2023.

New York City Department of Transportation (NYCDoT). (2021). Notice of Adoption relating to the demonstration or testing of motor vehicles with autonomous vehicle technology on any public highway in the city of New York. New York: New York City Department of Transportation, pp.1-7.

Polinori, T., Rice, T., & Monfalcone, M. (2021). Hazard analysis of concept heavy-truck platooning systems (Report No. DOT HS 813 065). National Highway Traffic Safety Administration. https://rosap.ntl.bts.gov/view/dot/55819/dot_55819_DS1.pdf. Accessed 14 March 2023.

Rene de Cotret, Y., Donga, J., Parent, M., Nathwani, R., Wilson, S., & Sterling, K. (n.d.). (rep.). Connected and autonomous vehicles in Ontario: Implications for the insurance industry. Deloitte, Province of Ontario. Retrieved from https://www2.deloitte.com/content/dam/Deloitte/ca/Documents/consulting/ca-EN-CVAV-Research-Insurance-Report-2018Apr25-Final-AODA.PDF. Accessed 14 March 2023.

Roberts, Jack. "Truck Platoons on the Horizon?", www.truckinginfo.com, 12 June 2018. Retrieved from https://www.truckinginfo.com/304199/truck-platoons-on-the-horizon. Accessed 14 March 2023.

SAE.org. 2021. J3016C: Taxonomy and Definitions for Terms Related to Driving Automation Systems for On-Road Motor Vehicles - SAE International. [online] Available at: https://www.sae.org/standards/content/j3016_202104/ Accessed 4 October 2021.

Sagalow, Z. (2021) Self-Driving Cars Raise Liability Questions for Insurers. https://www.govtech.com/fs/self-driving-cars-raise-liability-questions-for-insurers. Accessed 14 March 2023.

Schuelke Leech, B. A., Jordan, S. R., & Barry, B. (2019). Regulating autonomy: An assessment of policy language for highly automated vehicles. Review of Policy Research. https://doi.org/10.1111/ropr.12332. Accessed 14 March 2023.

Scottish Law Commission No. 258 (2022) Automated Vehicles: Joint Report, paragraph 2.22, https://www.lawcom.gov.uk/project/automated-vehicles/. Accessed 14 March 2023.

Senate Bill No. 996. 665b.

Taeihagh, A., & Lim, H. S. (2018). Governing Autonomous Vehicles: Emerging Responses for safety, liability, privacy, cybersecurity, and industry risks. Transport Reviews, 39(1), 103–128. https://doi.org/10.1080/01441647.2018.1494640. Accessed 14 March 2023.

Tennessee Code Annotated §55-30-106(a) (2019).

Thomasen, K. (2020) Out of their Cages and into the City: Robots, Regulation, and the Changing Nature of Public Space. https://commons.allard.ubc.ca/fac_pubs/633. Accessed 14 March 2023.

Tibken, Shara. "Waymo CEO: Autonomous cars won't ever be able to drive in all conditions." CNET.com, 13 November 2018. Retrieved from https://www.cnet.com/roadshow/news/alphabet-google-waymo-ceo-john-krafcik-autonomous-cars-wont-ever-be-able-to-drive-in-all-conditions/. Accessed 14 March 2023.

Ticoll, D. (2015). Driving changes: Automated vehicles in Toronto. Munk School of Global Affairs - University of Toronto. Munk School of Global Affairs, University of Toronto. Retrieved from https://munkschool.utoronto.ca/ipl/files/2016/03/Driving-Changes-Ticoll-2015.pdf. Accessed 14 March 2023.

Tkachuck, D., Dawson, D., Bovey, P., & MacDonald, M. (2018). (rep.). Driving Change: Technology and the Future of the Automated Vehicle. Standing Senate Committee on Transport and Communications, Senate Canada. Retrieved from https://sencanada.ca/content/sen/committee/421/TRCM/Reports/COM_RPT_TRCM_AutomatedVehicles_e.pdf. Accessed 14 March 2023.

U.S. Department of Transportation, Preparing for the future of Transportation: Automated Vehicles 3.0 (2018). Retrieved from https://www.transportation.gov/sites/dot.gov/files/docs/policy-initiatives/automated-vehicles/320711/preparing-future-transportation-automated-vehicle-30.pdf. Accessed 14 March 2023.

Vlacic, L., Wang, F, (2022) Let Us Reinstate the Public Trust in the Safety Benefits of Automated Vehicles. IEEE Intelligent Transportation Systems Magazine, March.

Walker Smith, Bryant. “Lawyers and Engineers Should Speak the Same Robot Language”, in Robot Law, ed. Ryan Calo, A. Michael Froomkin, and Ian Kerr. Elgar Publishing, 2016.

Wang, M. et al (2019). Benefits and Risks of Truck Platooning on Freeway Operations Near Entrance Ramp. Transportation Research Board. Retrieved from https://journals.sagepub.com/doi/pdf/10.1177/0361198119842821. Accessed 14 March 2023.

Wickert, G. (2021), Subrogating Automated Driving Systems and Autonomous Vehicle Failures, www.claimsjournal.com/columns/road-to-recovery/2021/12/28/307762.htm. Accessed 14 March 2023.

Wiggers, K. (2022). “Autonomous Cruise car encounter with police raises policy questions.” Retrieved from www.techcrunch.com/2022/04/11/autonomous-cruise-car-encounter-with-police-raises-policy-questions/. Accessed 15 February 2023.

Ziegler, M. S., & Trancik, J. E. (2021). Re-examining rates of lithium-ion battery technology improvement and cost decline. Energy & Environmental Science, 14(4), 1635-1651. Retrieved from https://pubs.rsc.org/en/content/articlelanding/2021/ee/d0ee02681f. Accessed 14 March 2023.

Zon, N., & Ditta, S. (2016). Robot, take the wheel: Public policy for automated vehicles. Mowat Centre, University of Toronto. Retrieved from https://munkschool.utoronto.ca/mowatcentre/wp-content/uploads/publications/118_robot_take_the_wheel.pdf. Accessed 14 March 2023.

16

Acknowledgements

The authors would like to acknowledge the kind assistance of the following individuals in the preparation of this white paper:

Adam Bell, Adriano Alessandrini, Alain Kornhauser, Amanda Kalhous, Andres Valencia Rodriguez, Andy Manahan, Angelo Lakrintis, Barbara Gray, Barrie Kirk, Bert Kaufman, Brad Templeton, Brendon Hemily, Brianna Puigmarti, Bryant Walker Smith, Carl Hansen, Carl Pucci, Charissa Iogna, Chelsey Colbert, Chirag Bhatt, Chris Plante, Corey Clothier, Craig Reid, Daniel Carballo, Darlene Anderson, David Crowley, Don Moore, Dylan Reid, Fahad Shuja, Geoff Wilkinson, Geoffrey Wood, Henry Wiercinski, Ignacio Tartavull, Isabella Foster, James Thompson, Jascha Franklin-Hodge, Jason Kuo, Joel Nascimento, John Niles, Jon Harrod Booth, Judy Farvolden, Julio Pereira, Katherine Sherif, Ken Vaughn, Khelil Khelil, Lauren Reid, Mark Fitzgerald, Martin Lord, Marty Collier, Matthew Roorda, Matthias Sweet, Michael Fenn, Michael Roschlau, Michael Sena, Michael Stewart, Michael Weisman, Mike Millian, Mike Ravenscroft, Murtaza Haider, Nadia Todorova, Nima Ashtari, Omar Choudhry, Paul Barter, Paul Kovacs, Peter Jones, Raed Kadri, Rich Steiner, Richard Mudge, Richard Soberman, Richard Yu, Chief Robert Martin, Robert Poole, Robin Chase, Roger C. Lanctot, Rohit Aggarwala, Ross McKenzie, Ryan Lanyon, Sam Saad, Sam Schwartz, Sandra Rothbard, Sarah Goldfeder, Scott Butler, Selika Josiah Talbott, Stephen Wickens, Subha Ramanathan, Sukhjit Dhaliwal, Ted Graham, Teresa di Felice, Thomas Barakat, Tim Lichti, Trevor Foster, Valerie Eisen, Veronica Eaton, Vignesh Ganapathy, Vince Mauceri, Wade Beebe, Will Schaefer, and Willa Ng.

17

Glossary

AI
Artificial Intelligence

ADAS
Advanced Driver-Assistance System

ADAS-Only
A vehicle with ADAS capability but no ADS capability

ADS
Automated Driving System

ADS-H
ADS that may be engaged or disengaged by the in-vehicle human driver or fallback driver, with the ADS able to override to prevent misuse

ADS-V
ADS engaged by the ADS itself or its teleoperator, such that an in-vehicle human driver or passenger cannot override

ADS-DV
ADS dedicated vehicle, such that there is no opportunity for user control, except emergency shutdown and system exit

AMR
Automated Mobile Robot; a term typically reserved for robots deployed in industrial, warehouse, or agricultural spaces where any human present has been trained in the robot's use

AV
Autonomous Vehicles or Automated Vehicle

CASE
Connected, Autonomous, Shared, and Electric

DARPA
Defense Advanced Research Projects Agency (affiliated with U.S. military)

DDT
Dynamic Driving Task

DTS
Draft Technical Standard (an ISO term)

HAV
Highly Automated Vehicles

HGV
Heavy Goods Vehicle

HOS
Hours of Service

ID
Identity

ISO
International Organisation for Standards

ISO METR
Proposed standard series: Management of Electronic Transport Regulations https://iso-tc204.github.io/iso24315p1/index.html

IT
Information Technology

LCV
Long Combination Vehicle

LIDAR
Laser Imaging Detection and Ranging

NCHRP
National Cooperative Highway Research Program

ODD
Operational Design Domain

OEM
Original Equipment Manufacturer

OTA
Over the Air

PDD
Personal Delivery Device (a type of sidewalk robot or PMR)

PMR
Public Mobile Robot; a term reserved for robots deployed in public, pedestrianized spaces (for example, a delivery robot)

SAE
Society of Automotive Engineers

SAE3016
The document SAE 3016-202104

Teleoperator
A connected and attentive human that has a degree of realtime oversight over one or more vehicles. This may include response to attending agents during an emergency, inspection, or enformencent event; it may also include a degree of latitudinal or lognitudinal control (steering, accelerating, braking). A teleoperator may thus be able to drive remotely, but this capability will always be extremely lmited, and often times not provided.

TR
Technical Reference (an ISO term)

TS
Technical Standard (an ISO term)

VIN
Vehicle Identification Number

WG
Working Group (an ISO term)

HATCH