



Driving Decisions

How Autonomous Vehicles Make Sense of the World

Sam Hind

palgrave
macmillan


Driving Decisions

Sam Hind

Driving Decisions

How Autonomous Vehicles Make
Sense of the World

palgrave
macmillan

Sam Hind 
Art History and Cultural Practices
University of Manchester
Oxford Road, Manchester, UK

ISBN 978-981-97-1748-4 ISBN 978-981-97-1749-1 (eBook)
<https://doi.org/10.1007/978-981-97-1749-1>

© The Editor(s) (if applicable) and The Author(s), under exclusive licence to Springer Nature Singapore Pte Ltd. 2024

This work is subject to copyright. All rights are solely and exclusively licensed by the Publisher, whether the whole or part of the material is concerned, specifically the rights of translation, reprinting, reuse of illustrations, recitation, broadcasting, reproduction on microfilms or in any other physical way, and transmission or information storage and retrieval, electronic adaptation, computer software, or by similar or dissimilar methodology now known or hereafter developed.

The use of general descriptive names, registered names, trademarks, service marks, etc. in this publication does not imply, even in the absence of a specific statement, that such names are exempt from the relevant protective laws and regulations and therefore free for general use.

The publisher, the authors and the editors are safe to assume that the advice and information in this book are believed to be true and accurate at the date of publication. Neither the publisher nor the authors or the editors give a warranty, expressed or implied, with respect to the material contained herein or for any errors or omissions that may have been made. The publisher remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

This Palgrave Macmillan imprint is published by the registered company Springer Nature Singapore Pte Ltd.

The registered company address is: 152 Beach Road, #21-01/04 Gateway East, Singapore 189721, Singapore

If disposing of this product, please recycle the paper.

Acknowledgements

Weirdly to me, writing this book was enjoyable. I was happy to leave my thesis behind, and the strange world of autonomous driving seemed too strange a world to ignore. The projects that lie behind this book have taken it in many directions, few of which I necessarily anticipated. The book itself was largely written in Cologne in the summer of 2022 and then again in Manchester through the rest of 2022 and—intermittently—throughout 2023. I never write as comfortably as I do when at home (wherever that may be) but have always enjoyed the feverishness of thinking and writing on the move. Accordingly, many of the arguments in the book came together whilst on the train as different parts of the Rhineland, the Siegerland, and the *Ostsee* zipped by.

Free from teaching in my first year at the University of Manchester allowed me to finish the project, for which I am grateful. I want to thank new colleagues in the Department of Art History and Cultural Practices (AHCP) for making the final stages of the book's development possible, especially the digital media and culture team—Claire Reddeman, Luca Scholz, Łukasz Szulc, and Shuaishuai Wang—who have made integrating back into UK university life easy.

I am grateful to colleagues at the Centre for Interdisciplinary Methodologies (CIM) at the University of Warwick, where I first presented the skeleton of this project in 2016. I want to thank Sybille Lammes for early encouragement and Noortje Marres for the invitation

to an interdisciplinary workshop on autonomous driving in December 2016.

Moving to Germany in 2017 was life-changing, and this book is largely the product of my time at the University of Siegen. I want to thank the GRK1769 Locating Media graduate school, where I was based for three years. In particular, Pablo Abend, Asher Boersma, Mine Gencel Bek, Hendrik Bender, Christoph Borbach, Magdalena Götz, Simon David Hirsbrunner, Timo Kaerlein, Max Krüger, Danny Lämmerhirt, Hannah Neumann, Eva-Maria Nyckel, Anna Lisa Ramella, Sebastian Randerath, Asli Telli, and Astrid Wiedmann. Sharing an office with Roger Norum was delightful, his attempts at various Yorkshire accents likewise. I'm not sure how to thank Tatjana Seitz—colleague, friend, co-author, housemate—but the hours spent discussing our respective projects were always inspiring. My life would have been a lot different without a shared love of the work of Phil Agre.

I also want to thank SFB1187 Media of Cooperation, the collaborative research centre where this book took shape. The pandemic made this period a little challenging, but it was the SFB that saw me through it. I want to thank Carolin Gerlitz and Tristan Thielmann for their organization of both the centre and the Navigation in Online/Offline Spaces project. Many—perhaps all—of the ideas in this book are the product of the conversations, discussions, debates, and events hosted by the SFB. In particular, I want to thank two people: Max Kanderske and Fernando van der Vlist. The countless video calls we scheduled throughout the pandemic were invaluable, and I am incredibly proud of the work we've done together. This book, again, would be very different without them.

To Alex Gekker, although it's been a while since we've seen each other in person, so much of the work in this book is indebted to our continued collaboration and friendship over the course of these last 12 (!) years.

In addition, I want to thank organizers and attendees of the following workshops and conference panels at which chapters of this book—in various stages of development—were presented: the Geomedia workshop (Siegen, Germany), Infrastructures of Sensemaking workshop (Konstanz, Germany), Speculative Ordinaries workshop (Lüneburg, Germany), Technology and Automation panel at Im/mobile Lives in Turbulent Times (Newcastle, UK), Human-Machine Learning panel at Re-situating

Learning: Making Sense of Data, Media and Dis/unities of Practice conference (Siegen, Germany), Artificial Intelligence panel at INDL-5 conference (Athens, Greece), Digital Twins, Cousins and Other Kins panel at STS Conference Graz (Graz, Austria), Sociotechnical Data Studies panel at NordicSTS (Oslo, Norway), and Where's the 'Intelligence' in AI? panel at STS Italia (Bologna, Italy). I count myself as incredibly privileged to have the opportunity to attend these kinds of events and am indebted to the warmth and kindness of all those willing to offer advice and support.

Thanks to my editor at Palgrave Macmillan, Marion Duval, and to production editor Raghupathy Kalyanaraman.

To friends and family in Manchester, Cologne, Yorkshire, and everywhere in between, I am always grateful. Thank you.

Funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) – Project-ID 190576466 – GRK1769 Locating Media.

Funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) – Project-ID 262513311 – SFB1187 Media of Cooperation.

Contents

1	Introduction: The Making of Decisions and Technological Decisionism	1
2	Mapping Decisions: The Promise of Mapless-ness	43
3	Training Decisions: Ground-Truthing the Interesting	67
4	Sensing Decisions: Perceiving, Classifying, Finessing	95
5	Demonstrating Decisions: Waymo's World	125
6	Securing Decisions: Sovereignty and Semiconductors	159
7	Relaxing Decisions: Making Driving Chill	193
8	Resisting Decisions: Coneheads in California	225
	Index	265

Abbreviations

ADAS	Advanced Driver-Assistance System
ADS	Automated Driving System
ALV	Autonomous Land Vehicle
AP	Average precision
AVPS	Autonomous Vehicle Passenger Service
CMU	Carnegie Mellon University
CPUC	California Public Utilities Commission
CVC	California Vehicle Code
CVPR	Computer Vision and Pattern Recognition conference
DARPA	Defense Advanced Research Projects Agency (US)
DMV	Department of Motor Vehicles (California)
E2E	End-to-end approach
FSD	Full Self-Driving
GPU	Graphics processing unit
GSP	General Simulation Program
KIT	Karlsruhe Institute of Technology
LADOT	Los Angeles Department of Transportation
mAP	Mean average precision
ML	Machine learning
MRC	Minimal Risk Condition
NHTSA	National Highway Traffic Safety Administration (US)
NTSB	National Transportation Safety Board (US)
ODD	Operational Design Domain

OR	Operational research
RO	Rider-only operations
SAE	Society of Automotive Engineers
SCI	Strategic Computing Initiative
SFCTA	San Francisco County Transportation Authority
SNV	<i>Stiftung Neue Verantwortung</i> (New Responsibility Foundation)
SSR	Safe Street Rebel
TFEU	Treaty on the Functioning of the European Union
TSMC	Taiwan Semiconductor Manufacturing Company
TTIC	Toyota Technological Institute at Chicago
VMТ	Vehicle Miles Travelled
VO	Vehicle operator
VRE	Vehicle Retrieval Event
WAD	Workshop on Autonomous Driving

List of Figures

Fig. 1.1	Diagram of machine learning's abductive decision-making. (Source: adapted from Amoores, 2022)	16
Fig. 1.2	Stanley processing 'pipeline'. (Source: Thrun et al., 2007)	24
Fig. 2.1	A timeline of automotive navigation and autonomous driving	47
Fig. 2.2	The 'brain' of an Uber self-driving vehicle. (Source: Uber ATG, 2019)	51
Fig. 2.3	'Map, perceive, predict and plan': The architecture of Waabi's MP3 model. (Source: Casas et al., 2021)	56
Fig. 3.1	Difference between coarse (top) and fine (bottom) annotations in the Cityscapes Datasets. (Source: Cityscapes, 2022)	74
Fig. 4.1	A simple diagram of perception, forecasting, and planning. (Source: Wellington, 2021)	106
Fig. 4.2	Joint perception and forecasting. (Source: Wellington, 2021)	107
Fig. 4.3	Examples of different 3D sensor data representations with 'freespace' visualized (right). (Source: Hu et al., 2020)	112
Fig. 4.4	A simplification of the streaming perception process. (Source: Li et al., 2020)	116
Fig. 4.5	Dynamic scheduling. (Source: Li et al., 2020)	117
Fig. 4.6	Processing video footage from an autonomous vehicle. (Source: Li et al., 2020)	119
Fig. 5.1	A simplified flow diagram of activities at an Acid Bessemer steel-making plant. (Source: Tocher, 1960)	131

Fig. 5.2	A visual depiction of the ‘field of safe travel’. (Source: Gibson & Crooks, 1938)	135
Fig. 5.3	A visual depiction of a path trajectory taking account of a ‘jaywalking’ pedestrian. (Source: Pek et al., 2020)	136
Fig. 5.4	Classification of Waymo-involved collisions. (Source: Schwall et al., 2020)	139
Fig. 5.5	Actual and simulated contact events involving pedestrians or cyclists. (Source: Schwall et al., 2020)	141
Fig. 5.6	An infographic of Waymo reaching one million ‘rider-only’ miles. (Source: Waymo., 2023)	145
Fig. 5.7	Waymo’s ‘conflict typology’. (Source: Victor et al., 2023)	147
Fig. 7.1	The ‘computer vision lesson not learned’ and the figure of the ‘cone guy’. (Source: Hotz, 2021)	201
Fig. 7.2	Latest release notes for Comma openpilot. (Source: Comma, 2023a)	206
Fig. 7.3	Calibration, Comma community style. (Source: Eyezenheim, 2022b)	207
Fig. 7.4	Calibration, behind the scenes. (Source: Eyezenheim, 2022a)	208
Fig. 7.5	Bug report. (Source: Comma, 2023b)	210
Fig. 7.6	Bug tracker. (Source: Comma, 2023c)	211
Fig. 7.7	Steering bug data. (Source: GitHub, 2022)	213
Fig. 7.8	Steering angle limit detected. (Source: GitHub, 2022)	214
Fig. 8.1	Minimal risk condition. (Source: SAE International, 2021)	240
Fig. 8.2	Cruise’s ‘in-app contextual cues’. (Source: Cruise, 2022)	244
Fig. 8.3	Frames from Safe Street Rebel’s TikTok video of how to ‘cone’ an autonomous vehicle. (Source: Safe Street Rebel, 2023c)	248
Fig. 8.4	‘Find a cone (they’re everywhere), then gently place it on the car’s hood’. (Source: Safe Street Rebel, 2023c)	249

List of Tables

Table 1.1	Key dates in the development of autonomous driving, 1985—2023	27
Table 3.1	A comparison between autonomous driving datasets	86



1

Introduction: The Making of Decisions and Technological Decisionism

This book is about the phenomenon of autonomous driving—not what it is from a strictly ‘technical’ perspective but where the dream of it has led us in recent years. It considers the ongoing, complex, costly, and contentious quest to automate driving offering a sustained focus on the ‘advance decisions’ deemed necessary to emulate the most ordinary of driving decisions, from turning a corner or merging onto a motorway, to stopping at traffic lights. This book considers how mapping, sensing, and algorithmic capabilities are gifted to autonomous vehicles through technical work performed by an array of actors in multiple locations: from computer science students performing data annotation work in industry-funded research centres to software developers building image-based machine learning (ML) models at autonomous vehicle start-ups.

The book therefore intends to complicate, and question, typical understandings of autonomous driving by challenging the technological determinism—and ‘decisionism’—that advocates offer of an inevitable, fully autonomous future. From the death of Elaine Herzberg in Tempe, Arizona, in 2018, to the suspension of Cruise’s operating licence in San Francisco in 2023, it is technological decisionism that best describes the operational ethics and logics of this mutating, nebulous, dream. The

book attempts to offer a fuller, and perhaps counter-intuitive, picture of the communities invested in this story.

One of those communities is the automotive industry, and whilst I will return to car manufacturers throughout the book—especially in Chap. 6—the industry, I argue here and throughout, has been both unwilling and unable to deliver such a future. Indeed, at various points throughout the last seven years, many have strived hard to prevent it from ever happening, whether through specific, strategic decisions, a lack of foresight and action, or the absence of certain kinds of skill and expertise. At the end of this introduction, I provide a bit more context as to why this is.

Instead, the book offers an understanding of autonomous driving from the perspective of a series of outside camps, each with their own motivation for bringing an autonomous future into being, and each with their own interest in seeing an end to the established automotive industry. In different ways and using different approaches, these outsiders have sought to cultivate a different sense of what autonomous driving might *mean* and *entail*.¹ What connects them, despite their many differences, is a commitment to machine vision and automated forms of sensing as bases on which machinic forms of decision-making can be built.

The book offers a plural account of decision-making practices in a sensor-driven, ML-dependent algorithmic age. It does so by drawing on seven years of research on autonomous vehicles. The book argues that to understand the concomitant shifts in how we think, move, and act with machines—in which all manner of sensors, algorithms, devices, and platforms are in action—they should be looked at through the lens of the decision, and how decisions are mapped, sensed, planned, simulated, secured, ‘relaxed’, executed, and—importantly—resisted. Decisions are themselves embedded within an array of human-machine practices, from the mapping of road data using ‘bolt-on’ consumer devices, to the calculation of road user trajectories in simulated software environments.

The book draws on a range of case studies from big tech companies (Waymo, Uber ATG), autonomous vehicle start-ups (Argo AI, Waabi),

¹ I use the terms ‘outsiders’ and ‘outside camps’ to refer to actors/groups who cannot be considered established automotive manufacturers but nonetheless have attempted to challenge the automotive industry through new technologies, strategies, and business models.

and university research initiatives (KITTI), to semiconductor chip manufacturers (Taiwan Semiconductor Manufacturing Company [TSMC]), open-source projects (Comma), and anti-autonomous vehicle activist groups (Safe Street Rebel) all attempting to ‘disrupt’ the automotive industry in different ways. Thus, the book is hopefully a route through these various communities and their motivations, who barely agree on what autonomous driving means, let alone how to deliver it.

A great many of the technical processes and practices devised and deployed to help map, train, sense, simulate, and execute the decision-making abilities of an autonomous vehicle or its underlying systems ordinarily constitute ‘boundary objects’ (Leigh Star & Griesemer, 1989; Leigh Star, 2010) in the most straightforward of senses—material entities offering interpretive flexibility to multiple, interconnecting groups working more-or-less collaboratively. Such entities—from ML models to bolt-on driver-assist devices—are very much ‘the stuff of action’ (Leigh Star, 2010, p. 603), more-or-less binding technical workers together.

Yet the purpose to which these material entities are put is less clear—despite, in the most nominal of senses, being destined to help deliver the automation of driving. The ‘boundary work’ is fastidious, committed, and, as my outlining of the eras or ‘phases’ of autonomous driving later in the introduction suggests, wholeheartedly incremental. Yet so often and throughout these various communities and settings, the goal of such collaborative work is either wholly undetermined or wildly different from those in these other settings, equally committed to delivering autonomous driving.

Rather than considering this a threat to the emancipation of human drivers trapped in boring commutes, locked into big tech-driven economic models, or the doubtful delivery of safer roads or fewer automotive accidents, the book offers an insight into the state of technological decisionism in the contemporary world: powered by intersecting beliefs about the societal need to automate driving and how it might be done.

The remainder of this introduction is divided into three sections—decisions, operations, phases—and a chapter outline.

Decisions

Decision-making in a computational context has a long history, especially in respect to machine vision and control. In recent times, Bader and Kaiser (2019) have discussed the organizational nature of algorithmic decision-making and the role of interfaces in brokering human access to algorithmic decisions. Dencik and Stevens (2021) discuss the role of data-driven hiring systems in selecting job applications, designed to streamline decision-making for administrators and managers. Louise Amoore (2020) discusses the ethics of algorithmic decision-making, especially with regard to the ‘weightings’ within ML processes and the need to place moments of decisions ‘beyond doubt’ (Amoore, 2020, p. 147). For the purposes of this book, however, there are two trajectories worth outlining. Firstly, statistical ‘pattern recognition’ work beginning in the 1960s, laying the foundations for subsequent machine vision/ML work. Secondly, operational research (OR) and cybernetics work beginning in the late 1950s, which revolutionized ideas around automated forms of decision-making.

A History of Decisions

Statistical decision theory emerged as a new field of statistics in the late 1930s, courtesy of Abraham Wald, and wartime operations research in the US (Mendon-Plasek, 2020). As Mendon-Plasek explains in his pre-history of machine learning, decision theory made its way into early Optical Character Recognition (OCR) pattern recognition work, a key constituent in ongoing machine vision research.² As he writes, ‘decision theory offered a ready-made toolkit that could compare engineered pattern recognition systems in more meaningful ways using the concept of a “loss” function’ (Mendon-Plasek, 2020, p. 49), through the integration of so-called weights indicating accumulated loss (i.e. errors, correct

² CVPR, a significant annual academic conference in the world of computer vision, stands for ‘computer vision and pattern recognition’. Work on computer/machine vision and pattern recognition, thus, can be considered as disciplinary bedfellows.

judgements) at each step along the decision-making process. More generally, Mendon-Plasek argues that:

Contemporary debates about the generalizability of machine learning in social decisions rehearses many of the same debates pattern recognition researchers had with each other in the 1960s about how to compare different learning machines using different data sets. (Mendon-Plasek, 2020, p. 57)

In short, that rather than the 1960s constituting an oft-considered ‘AI winter’ in which progress towards certain AI goals had stalled, work in pattern recognition can be seen as a precursor to later progress in machine learning itself. Especially, that is, with regard to the quest for generalizability, and generalizable methods that could be applied across institutional contexts. Here, work on pattern recognition had contributed to the later, subsequent development of machine approaches to decision-making. Refracted through this earlier work in pattern recognition, statistical decision theory found new image-based applications. By making statistical use of prior observations in pattern recognition work, OCR researchers, as Mendon-Plasek (2020) argues, had all but invented supervised learning—a now-common ML and machine vision technique.

What this application of decision theory to OCR research did was enable the specificity of OCR work in each institutional domain it was then being applied in (i.e. in the postal service, education, US military) to be consolidated. In short, that such pattern recognition methods could offer a form of ‘generalization ability’ (Alpaydin, 2016, p. 40; Steinhoff & Hind, 2024) they were not previously capable of, with previously context-specific classification criteria (what Mendon-Plasek [2020, p. 48] calls ‘antithetical feature values’) now made comparable through the loss function. In work by Burroughs Corporation researcher CK Chow, a significant US computer company through the post-WWII era, the ‘optimal possible performance for any system’ could now be compared with another ‘according to a particular decision criteria’ (Mendon-Plasek, 2020, p. 50), as well as being able to pinpoint degradation issues associated with pattern recognition systems. Put simply, this is an early example of the ‘logic of domains’ (Ribes et al., 2019) that has ensured the

formalization of AI as built on the implementation of generalizable methods to other applied areas or ‘domains’, including the development of autonomous vehicles with computer vision/pattern recognition capabilities.

Operational research (OR) and management cybernetics are also key antecedents to contemporary work on machine vision and machine learning (Pasquinelli, 2023). Stafford Beer’s *Cybernetics and Management* (Beer, 1959) and *Decision and Control* (Beer, 1995 [1966]) can be seen as foundational texts of such thinking. Both are applications of earlier work, such as Norbert Wiener’s *Cybernetics* (Wiener, 1973 [1948]), to contemporary business contexts, where computational systems were beginning to be used within firms for an array of different tasks, from stock-counting to payroll administration. What Beer envisioned was integrated systems for the automated making of decisions. OR was considered a scientific practice that could be applied to the business of contemporary firms and to the management of their various activities. More than this, it could help firms—and specifically their executives—make better decisions. As Beer wrote at the time, ‘the whole idea of using hard science as an intrinsic part of the managerial process is alien to many’ (Beer, 1995 [1966], p. 6), and thus, Beer and his contemporaries within the fledgling world of OR and management cybernetics were seen as outsiders: those with specific, perhaps even esoteric, scientific, or philosophical knowledge, but otherwise lacking business acumen or familiarity with the day-to-day decision-making of the firm.

As Beer considers, for individual managers, ‘insight, value judgment, flair, acumen, maturity and experience *count*’ (Beer, 1995 [1966], p. 6, authors’ emphasis). For firms with management teams, ‘the climate of opinion, fashion, reputation, maintenance of face, dominance and every kind of personal relationship also count’ (Beer, 1995 [1966], p. 6). Thus, as Beer suggests, replacing all these dimensions and dynamics of ‘frail humans’ with ‘infallible electronic computers’ (Beer, 1995 [1966], p. 6) was neither especially possible nor desirable.

Yet, Beer presented a situation—ostensibly one for a manager or management team—that depended on the making of a key decision. In this case, a decision is about a particular industrial plant or site with a collection of tools and pieces of necessary equipment within it. In a simple

illustration of this situation, a manager might be tasked with drawing on the thoughts and expertise of relevant parties within the firm, from the engineers to accountants. The manager's task, as Beer saw it, 'with a decision to reach' involved 'rolling up all such small and isolated verdicts into a ball to produce a *consolidated verdict* about the relative merits of *A* and *B* "on the whole", "in the long run"' (Beer, 1995 [1966], p. 7, emphasis added). In making this 'consolidated verdict', managers would require 'plenty of judgement' alongside 'the weighing of evidence' such that 'it is almost a juridical proceeding' (Beer, 1995 [1966], p. 7). This aspect of decision-making—of weighing and of weights—can be seen as integral to all kinds of decisions, whether human or computational.

Beer thus operationalized this dimension: 'as situations become more complicated, so the problem of establishing an adequate scientific basis for decisions becomes greater' (Beer, 1995 [1966], p. 8). Yet in understanding that ideal scientific conditions do not exist within such an economic, organizational context, the manager 'must use whatever information there is in an attempt to establish the *probabilities* that one [decision] is better than another' (Beer, 1995 [1966], p. 8, authors' emphasis). Here, the conditions of the making of the decision—what Beer refers to elsewhere as the 'environments of decision' (Beer, 1981 [1972], p. 181)—are integral to how one reaches any such decision, by applying necessary weight to each implicated factor.

Further, that in evaluating the situation at hand, it may be necessary to look beyond it: 'if there is nothing *within* the *A/B* situation to indicate which is the better choice, then we should look *outside* the situation' (Beer, 1995 [1966], p. 8, authors' emphasis). Here, context matters, and the wider environment of the decision should always be brought to bear on the making of the decision itself, even if 'exterior' to the situation at hand. Pithily expressed, decision-making is as much about *not* making a decision as making a decision, Beer tells his management audience.

The point Beer makes is that computers can help managers make 'better' decisions, by performing this weighting and assessment of criteria more efficiently, and at scale, compared to managers on their own. Rather than OR or management cybernetics being considered a threat, Beer positions them as useful to their evolution into *better managers* in a computational age—despite resistance. In this, Beer establishes the key tenets

of organizational decisions and decision-making (factors, information, weights/probabilities, situations, judgements/verdicts), and that managers in firms—alongside many more other people, as this book will detail—have their own beliefs about the capabilities of computational systems to make decisions either for them or instead of them.

Theories of Decision-Making

In light of these trajectories, I want to consider what a theory of decisions and decision-making might now look like—not from a statistical or managerial standpoint, but from a machine vision perspective. Theoretically, such an approach would be indebted to German media scholar Florian Sprenger, whose work on ‘micro-decisions’ (Sprenger, 2015, 2020, 2021) can be considered foundational. Building on Alexander Galloway’s work on internet protocols (Galloway, 2004), Sprenger considers micro-decisions as occurring at the ‘level of technical infrastructures’ (Sprenger, 2015, p. 20), constituting the ‘nodes of networks’ (Sprenger, 2015, p. 20) like the internet itself, that function as tiny interruptions managing the flow of data through such networks. As Sprenger contends:

In terms of the logic of decision-making, the basis of all computers and their networks is not only structured on the level of binary code but also on that of the protocols that produce connections and disconnections, participation and non-participation. (Sprenger, 2015, p. 21)

Micro-decisions consequently underpin computation as its networked form is known and understood today. The key differences I want to make, however, are two-fold.

Firstly, that the approach in the book does not exclusively work with the notion of temporality and the temporal dynamic of decisions, nor that there is a specific interest in the ‘micro-’ decision occurring beneath the level of human attention and response. As Sprenger more recently writes in relation to the autonomous vehicle, the temporal form of the micro-decision ‘is an effect of the relation between the sheer number of

calculations and the velocity of automated processing’ (Sprenger, 2021, p. 160) in any given situation. The ‘intervals’ of the micro-decision, thus, ‘are too short for human attention’ (Sprenger, 2021, p. 106). Here, whilst the book draws at times on a similar literature around the temporality of decision-making, via Katharine Hayles’ idea of ‘cognitive assemblages’ (Hayles, 2017; Hind, 2022), like Sprenger (2021) himself, the book does not exclusively focus on temporality as a quality of a decision.

But secondly, that the approach centres the *making* of decisions rather than the decision, per se. In this, decisions are always ‘made’, usually through entrenched, repeated, encultured, sedimented actions that can be understood as specific kinds of decision-making practices, necessarily with their own dynamics, features, and operative components. For instance, in how remote ‘mobility managers’ might be imagined to be able to intervene in the decision-making of autonomous vehicles (Hind, 2022). Indeed, that they can both be *specified* and *located* within the kinds of case studies the book considers, as well as classified to some degree based on what the decisions are designed to enact and how they are executed. As Sprenger (2021, p. 170) considers in relation to machines, ‘we need a conceptual framework [of decisions] that helps us understand their mode of power’.

Considering the notion of the decision a little more, Andrew Dwyer suggests algorithms don’t, or can’t, make decisions which are instead a ‘distinctly human practice’ (Dwyer, 2020, p. n.p.). In this, Dwyer understands calculations as different from the making of decisions, following Derrida (1992), in which ‘calculation can form part of a decision, but is not equivalent to it’ (Dwyer, 2020, p. n.p.).

In this, it is best to be clearer about what kinds of decisions—and specifically machine vision, ML-based decisions—are being discussed. In the book I am primarily interested in decisions that do not, necessarily or always, surface above the technical system in question. In these instances, decisions (strictly only calculations, in Dwyer’s formulation) are not always made by human operators at the moment of execution. Whilst they may be *programmed* or otherwise designed by engineers, developers, or other technical workers in advance to make such decisions (and as such satisfy Dwyer’s definition above), at their literal, situational execution—say at a road junction or at the point of the merger of two

lanes—such decisions are made by the relevant mapping, sensing, planning, and control modules of the autonomous vehicle in question. Here, I would argue that these decisions are executed by a machine, *as well as* ordinarily made by engineers, developers, or other such technical workers responsible for designing, training, testing, and running the machine.

Most importantly, following Mackenzie (2017) and Roberge and Castelle (2020) decision-making in an ML context can be understood as *navigational*, in that it extrapolates and projects itself into future states. Within an autonomous vehicle context, this is most evident in the steps along the ML ‘pipeline’ (see later) that come after perception and object-recognition, namely ‘forecasting’ and ‘motion planning’. In these navigational steps, the focus of Chap. 5, ML models must learn to accurately predict the future state of actors, including the ‘ego vehicle’ they support. In this a model must contend with a vast volume of possible trajectories and future states, according to parameters set by ML practitioners. Thus, decision-making, or the execution of decisions, does not end in a single, final outcome: ‘the decision’. Instead, decision-making must always work along a *navigational path*, finding the best way to proceed, having taken account of the possibly converging paths of other actors in the world. This is what Roberge and Castelle (2020, p. 13) refer to as machine learning’s ‘quest for agency’, that ‘machine learning models are readymade as (semi-)autonomous’ in which the ‘act of classification, whose accuracy is optimized during training, can become an act of decision-making during deployment’ (Roberge & Castelle, 2020, p. 13). In other words, that ‘machine learning culture is more directly involved with the possibility of taking *action*’ (Roberge & Castelle, 2020, p. 13, authors’ emphasis) as compared with other algorithmic techniques.

Roberge and Castelle (2020) use Latour (1986) to refer to the ‘cascade’ of decision-making moments or instances, whilst one could also talk of their *concatenation*, in that decision-making moments generate the spark for subsequent concatenated, decision-making moments. Rather than spontaneous or organically initiated moments, however, ML models rely on calculable pathways or trajectories along which decisions actualize. Virtual environments, naturally, are crucial agents in such work, providing the space for such concatenations to occur, but also the space for

trajectories to be optimized, and re-optimized, played, and re-played (with).

Florian Jatón (2021, p. 23), following the work of French cognitive scientist Jacques Theureau (2003), talks similarly of ‘courses of action’. Here, Jatón focuses on the ‘accountable chronological sequences of gestures, looks, speeches, movements, and interactions between humans and nonhumans whose articulations may end up producing *something*’ (Jatón, 2021, p. 23, authors’ emphasis). Part of Jatón’s reasoning for using a more ‘generic definition’ (Jatón, 2021, p. 23) to consider algorithmic work is that it helps ‘resist the supposed abstraction of computer science work’ (Jatón, 2021, p. 23) in which proponents are prone to mystifying rather than clarifying the work undertaken, playing up its novelty rather than establishing similarities or lineages with prior or parallel work.

Offering examples of courses of action that ‘produce something’ Jatón mentions ‘a piece of steel, a plank, a court decision, an algorithm’ (Jatón, 2021, p. 23), however with cascading decision-making moments, or a concatenation of such events, these things likely result in further *enactments* rather than discrete or boundable objects. In this, it might even be unwise to draw on Jatón’s definition of activities as sets of ‘intertwining courses of actions sharing common finalities’ (Jatón, 2021, p. 23) if reference to ‘finalities’ is objectionable, even if Jatón’s decision to foreground doing (ground-truthing, programming, etc.) rather than being (ground-truth, programme) echoes the move I make in this book from decisions to decision-making.

Decisions: Advance, Deferred, Distributed, Discretionary

This is primarily why I am interested in ‘decision-making’ rather than decisions, per se. The term has a manifold connotation that considers both a decision as *made* (produced, manufactured) as well as *executed* (enacted, completed). In this the human work required to gift an autonomous vehicle decision-making capacities as well as the *machinic capability* to enact those decision-making capacities at any given moment are both acknowledged. In this, I talk expressly about ‘advance decisions’ made in

advance of executed driving decisions (turning, merging, stopping), in order for these decisions to be made—or, indeed, to not be made. I talk in Chap. 5 about how decisions are often ‘deferred’, rolled into continuous, iterative, feedback loops of decision-making.

For Sun-ha Hong ‘predictive systems, by definition, reconfigure the existing distribution of decision-making relations’, asking ‘which parts of the decision are considered sufficiently unknowable that they are open to (unequal) negotiation’ (Hong, 2022a, p. 6)? Likewise, that ‘prediction serve to *reallocate* discretionary power across different actors, and additionally to *obfuscate* the continuing role of discretionary power in decision-making’ (Hong, 2022a, p. 6, authors’ emphasis). For Hong, ‘discretion ... describes the always unequal distribution of the power to *define the situation*’ (Hong, 2022a, p. 6, authors’ emphasis), the question, however, is how or whether discretion works at all in ‘closed-loop’ decision-making, where algorithmic outputs *aren’t* presented to discretionary actors but continuously fed back into and fuelling more decision-making. Where, if at all, does this discretionary power lie? As Hong (2022a, p. 8) reiterates:

Prediction grammatises – renders flexibly replicable, habituates, provides a template for – a widespread extraction of discretionary power: the spaces of practical ambiguity, the gap between rule and case, the moments of situational judgment, that were always unequally distributed across different subjects in the first place. (Hong, 2022a, p. 8)

Alongside this question of how decision-making relations are distributed, John Law also writes that ‘decision making ... may be understood as the performance of certain forms of overlapping distribution’ (Law, 2002, p. 146). In this, the distribution of decision-making is greater than the allocation of decisions as if they were goods or objects to hand out. Here, ‘distributions resemble one another or may at any rate ... be *made* to resemble one another’ (Law, 2002, p. 146, authors’ emphasis). Talking specifically about political decisions rather than technological ones, per se, Law considers how decision-making assumes strategic importance and through which some form of masking or camouflage allows certain decisions to blend into one another, or present as singular instead of multiple.

Accordingly, decision-making practices can be understood as ordinarily involving the enaction, or deployment, of various kinds of tricks and tactics that resist an instrumentalized reading of the pure, rational, computational execution of decisions. In other words, that there are other things going on when decision-making takes place, beyond, although routinely in service to, the strict making of decisions.

One other aspect of the distribution of decision-making that Law considers is how “options” are brought into being’ (Law, 2002, p. 147), distributing ‘that which is possible, and that which is not’ (Law, 2002, p. 147). In technical terms—something that will be considered at various points in this book—the phrase ‘options’ might be substituted for ‘parameters’. Thus, parameters—settings that establish the operational limits of certain kinds of simulations, for example—distribute but also encode, ‘that which is possible, and that which is not’ (Law, 2002, p. 147). Or more precisely, drawing on the relevant history of pattern recognition: *that which is statistically likely, and that which is not*. Yet perhaps unlike other possible options not considered that are ‘removed from the universe of possibilities’ or indeed might not even have been ‘conceived as options in the first place’ (Law, 2002, p. 147), the setting of technical parameters can be understood as theoretically endlessly modifiable.

Here, parameters play the role of setting the limits of the conditions to be tested. Where there may be greater similarities is in the practice of setting those limits, of what to include and exclude, what to studiously ignore, or unknowingly discount. Everywhere in the development, testing, and deployment of autonomous vehicles exclusionary practices can be found, and everywhere evidence of the actualization of the restrictive qualities of parametrization can be witnessed. Although sometimes arbitrary, such decisions are routinely driven by computational, and thus political, constraints: the hard limits of mobile hardware carried within an autonomous vehicle, as well as the reputational limits of vehicle safety data carefully released to the general public. This is broadly what Ludovico Rella understands as the ‘material political economy of the epistemology of computation’ (Rella, 2023, p. 1).

Decisionism

I want to refer to the proliferation of machine vision/ML-dependent decision-making as ‘technological decisionism’ (Parisi, 2017). Technological decisionism refers to how decisions are valued only for their determination of clarity rather than correctness, what Bessner and Guilhot (2019) call, in broader terms, a ‘decisionist imagination’.

Firstly, that in order to enact and enable control over forms of life, one must first fabricate, and then encode, decisions within the fabric. In this, social, cultural, political, or technological processes that may not have required decisions, nor the making of decisions, are reformatted through decision-making architectures imposed upon them. Secondly, that this encoding of decisions must necessarily be delivered and buttressed by autonomous systems, such that one is able to ensure the making of decisions is carried out within an operational vacuum (e.g. the AI ‘pipeline’ without any exhaust or outlets), free of bias, free of motive, and free of alternative ideas and ideologies about the form of life subject to decisionism, in which the setting, management, and making of decisions follow a prescribed path.

For Sun-ha Hong, this maps onto how predictive policing tools and similar algorithmic decision-making systems ‘institute a particular *decision* about how futures must work’ (Hong, 2022b, p. 384, authors’ emphasis) and how future life must be constructed. The power of these efforts is not in ‘proving’ the validity of decisionism as a worldview or orientation, as Hong suggests, but in demonstrating it ‘by offering shareable and visceral experiences that [are designed to] impart a different way of seeing’ (Hong, 2022b, p. 386). This book is, inter alia, about such demonstrations of technological decisionism, as they abound in the world: from the responsibility of Waymo vehicles, to the ‘vibes’ sought by Comma users.

As Parisi continues:

Machine-learning algorithms do not simply perform nonconscious patterns of cognition about data, exposing the gaps in totalizing rational systems, but rather seem to establish *new chains of reasoning* that draw from

the minute variations of data content to establish a machine-determined meaning of their use. (Parisi, 2017, p. n.p., emphasis added)

Here, returning to Roberge and Castelle (2020), the power of the chains, cascades, and concatenations of decision-making, and the enforcing of such reasoning upon different domains and lifeworlds are fundamental to understand. Machine learning's recursive 'quest for agency' (Roberge & Castelle, 2020, p. 13) leads it to imposing—or constantly seeking to impose—itsself on situations otherwise nominally open to alternative logics and forms of reasoning.

However, I diverge from Parisi (2017) in how the apparent 'clarity' of technological decisionism has given way to a splintered sense of decision-making, whereby clarity is neither sought nor secured but necessarily masked. In other words, that technological decisionism needn't always search for clarity, or at least surface or evidence such clarity or finality. Returning to Law (2002), machine learning routinely defers and distributes the making of decisions, such that whilst clarity might be sought, the search for it is hidden or blurred in some manner. This is not to say that 'indecision' is favourable—certainly not when an autonomous vehicle on a public road blocks a junction—but that the deferral of the decision is deemed desirable in some sense. No decision is ever final, but that does not necessarily make it contestable. As the book will hopefully show, precisely because each decision-making moment is provisionally assured, it is provisionally able to suspend the prospect of contestation indefinitely.

Louise Amoore similarly examines the rise of ML-driven decision-making (Amoore, 2022). Here, Amoore considers first how 'rules-based computation and decision was critical to the formation of postwar international liberal order' (Amoore, 2022, p. 26), asking 'what happens when the machine learning function displaces it?' (Amoore, 2022, p. 26). More specifically, Amoore is concerned with how ML 'functions' become the mechanisms through which political orders are now generated, as opposed to rules-based algorithmic processes, dominant before the emergence of machine learning. What Parisi calls technological decisionism is reframed, or refocused down a level, to the functional operation of any ML system itself, as well as the relationship(s) formed between such systems and the political entities or bodies increasingly dependent on them for

decision-making. As Amoore concludes, ‘in a sense, machine learning’s *raison d’être* is to generate outputs that are in excess of the formulation of rules’ (Amoore, 2022, p. 26), operating in a space where the ontological order and arrangement of objects in the world are governed by an unruly, calculative process.

In this, Amoore also suggests that what makes machine learning different is that it proceeds from (desired) target to input(s) and layers, rather than from input and rules to output, ‘abductively working back to adjust the parameters of [a] model in order to converge on the target’ (Amoore, 2022, p. 28). This is what she refers to as ‘retroactive design’ (Amoore, 2022, p. 28). It is this inversion of the ordinary decision-making process, in which the ‘end target’ becomes the ‘starting point’ (Amoore, 2022, p. 28), that ‘significantly reconfigures the relationship between the formulation of a political problem and the proposition of a solution’ (Amoore, 2022, p. 28). It is also a key feature of technological decisionism, in which the reaching of a decision takes precedent (Fig. 1.1).

This technological decisionism, however, can also be seen as a modulation of decentralized computational networks, what Matteo Pasquinelli (2021), through a reading of Friedrich von Hayek’s theory of markets, calls ‘mercantile connectionism’ (Pasquinelli, 2021, p. 161). Here, he understands Hayek’s theory of market activity as involving the acquisition of knowledge ‘through the act of classification or pattern recognition’ (Pasquinelli, 2021, p. 162) before, neurologically, ‘a topology of

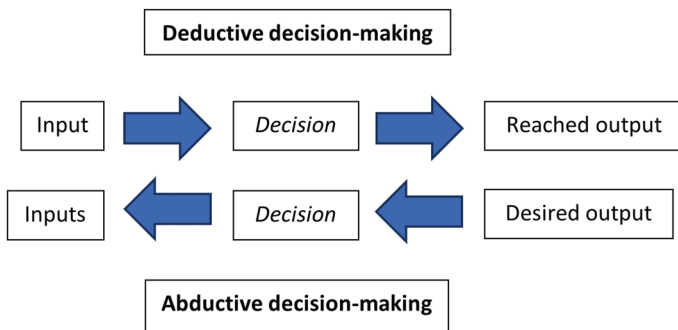


Fig. 1.1 Diagram of machine learning’s abductive decision-making. (Source: adapted from Amoore, 2022)

connections to take decisions’ (Pasquinelli, 2021, p. 162) is determined. As Pasquinelli contends, ‘eventually, in Hayek’s political intention, connectionism and neural networks provide a relativist paradigm to justify the “methodological individualism” of neoliberalism’ (Pasquinelli, 2021, p. 162) itself. Connectionism, thusly—the foundational theory of contemporary machine learning—likewise undergirds and provides a neurological, ‘natural’ justification for neoliberal market activity.

Amoore (2022) also considers the art and act of decision-making in a slightly different register, returning to the allure of technological decisionism. In this Amoore explains how algorithms are considered to work rationally and logically, expelling any ‘unreason’ (Amoore, 2020, p. 110) from the systems through which they operate. Yet as Amoore contends, contemporary ‘accounts of moral panic amid the madness of algorithms’ (Amoore, 2020, p. 110) and their unintended effects, forget that unreason is central to computational logic itself, as well as the ‘capricious incalculability within our twenty-first-century modes of algorithmic decision’ (Amoore, 2020, p. 110). In other words, that far from acting sensibly and rationally, algorithms—and especially ML models—incorporate all kinds of ‘madness’. Ostensibly, like the madness of markets themselves, as Pasquinelli’s reading of Hayek’s connectivism might have it.

As Amoore eloquently puts it, such algorithmic ‘decisions are mad because they can never fully know the consequences and effects of their own making’ (Amoore, 2020, p. 112), operating blissfully, yet logically unaware of the effects their calculative actions have on, and in, the world itself. It is the algorithm-as-aberration, or the moral panic of algorithms, that is frequently invoked in the world of autonomous driving, each time in doing so serving to erroneously demarcate the sensible, reasoned decision-making of particular ML models-in-action and those that have erred, or strayed, from their rational foundations.

As Chap. 5 argues, considerable energy is dedicated to testing the thresholds at which ML algorithms begin to act a little too imaginatively, or carelessly, for the chosen task. What this book hopefully offers, therefore, is an account of how certain people—engineers, programmers, software developers, and the like—have applied the kind of moral algorithmic judgement that Amoore unpacks. In short, to examine the role of the algorithmic judgements made by practitioners as they work on, and with,

autonomous vehicles, and to offer ‘an insight into how the algorithm enrolls and deploys ideas of unreason to function and to act’ (Amoore, 2020, p. 117). As Amoore sums it up: ‘the unreason and the excess are not aberrations at all’ for the ML models they power, ‘but are the condition of possibility of action’ (Amoore, 2020, p. 119). What space is there, or indeed, what appetite is there, for a level and kind of algorithmic experimentation that risks the making of such unreasonable decisions not just in public but on the social road? What price is paid for the ML model’s errant sense of adventure, as Amoore (2020) would have it? What do the decision-makers decide? Conversely, following Andrejevic (2019), who decides the decision-makers?

Operations

Methodologically, the book offers an ‘operational analysis’ of the phenomenon of autonomous driving. This approach centres on a critical evaluation of the technical operation of autonomous vehicles and the work underpinning this operation. Drawing on, and synthesizing, different methodological approaches—technography, operational analysis, the study of situations—the book is indebted to the work of others who have considered the construction and calibration of technological operations in innovative, and inventive, ways.

Technography and Techniques

Taina Bucher’s (2018) use of technography aids, for instance, the ‘mapping [of] the operational logics of algorithms’ (Bucher, 2018, p. 60). Technography, for Bucher, ‘is a way of describing and observing the workings of technology in order to examine the interplay between a diverse set of actors (both human and nonhuman)’ (Bucher, 2018, p. 60). Here, Bucher is not alone in offering a definition of technography. As she states, Grant Kien has previously described it as an ‘ethnography of technology’ (Kien, 2008), whilst Steve Woolgar suggested it offered a way to ‘tease out the congealed social relations embodied within technology’

(Woolgar, 1998, p. 444) requiring, in his words, ‘sustained empirical study in technical settings’ (Woolgar, 1998, p. 444).

In the French tradition François Sigaut considers technography as an ‘anthropology of technics’ (Sigaut, 1993, p. 422) and as a method to establish ‘technical facts’ (Sigaut, 1993, p. 424) through the study of ‘technical acts’ or operations where an operation is ‘someone doing something’ engendering ‘material change’ that can be ‘usefully observed’ (Sigaut, 1993, p. 425). As Sigaut continues, such technical acts or operations typically occur ‘as parts of a sequence that can be called a “path”’ (Sigaut, 1993, p. 423) such as the brewing of beer, the making of shoes, or the processing of cereals. In the case of brewing beer, such paths, as Sigaut explains, include saccharification and fermentation, which cannot ordinarily be skipped or avoided. As Sigaut continues:

In Europe, the brewing of beer is preceded by growing barley and hops and culturing yeast [other paths]; it demands a variety of devices that have had to be made by the corresponding craftsmen; it burns fuel, and so on. Step by step we realize that all the paths present in any one society are interwoven, in some way or other, into a sort of *network*, which is in fact the economic organization of that society. (Sigaut, 1993, p. 426, authors’ emphasis)

Similar to Bucher, Bernhard Rieder offers a study of ‘algorithmic techniques’ (Rieder, 2020) such as machine learning, through the Simondonian notion of ‘mechanology’ (Simondon, 2017). In co-authored work, a synthetic approach to studying ‘algorithmic systems’ (Rieder & Skop, 2021, p. 4) is offered, indebted both to Bucher and Rieder’s own work, in order to account for the distributed operation of ‘machine moderation’ (Rieder & Skop, 2021, p. 2) on the web. Van der Vlist et al. likewise employ a technographic approach to study the emergence of ‘big AI’, enabling a ‘detailed examination of the material aspects of technology by directly reading various publicly available documents generated by and related to technical systems’ (Van der Vlist et al., 2024, p. 4). For Hind et al. (2024), technography offers the opportunity to study the role of ‘challenges’ in catalysing the development of AI products and systems.

Operations and Operational Analyses

Whilst the origins of the study of computational operations are traceable back to OR in both military (1940s) and business (1950s–) contexts, only recently have scholars returned to the question of operations and how to study them. The impulse has evidently been the rise of AI, machine learning, and the automation of the production of images (Pasquinelli, 2023). For media scholars long versed in various kinds of representational, semiotic, and visual/textual analyses, this returning interest has stimulated new theoretical and methodological approaches for understanding the ‘politics’ of automated operations (Mezzadra & Neilson, 2019).

Kathrin Friedrich and AS Aurora Hoel deploy an operational analysis to ‘systematically observe and critically analyze ... [the] situated, interventional and multilayered entanglements’ (Friedrich & Hoel, 2023, p. 50) of digital media applications such as clinical imaging systems. As they contend, ‘media researchers are increasingly faced with shifting entanglements between physical and virtual layers of operations and interventions – with systems where human and non-human agencies intertwine and intra-act in often seemingly inscrutable ways’ (Friedrich & Hoel, 2023, p. 51). In this, media researchers are necessarily forced to understand how typical media research objects (images, video frames) are enrolled into automated processes, ordinarily hidden from view.

To understand these processes better, Friedrich and Hoel consider the media integral to the operation of such systems as ‘adaptive mediators’ (Friedrich & Hoel, 2023, p. 51). Here, Friedrich and Hoel build on substantial work into the ‘operational’ nature of media and, specifically, an operational reading of images (Farocki, 2004) and data (Walker Rettberg, 2020). As they further state, following Simondon (2017), ‘[i]n their role as adaptive mediators ... technical objects enact a new resolution of the human-world system, a material reconfiguration of the world that releases new potentials for perception and action’ (Friedrich & Hoel, 2023, p. 56). In their function as adaptive mediators, operational media act as engines within the world, offering enactive possibilities in near-infinite circumstances.

For Mezzadra and Neilson, there is an interest in how the ‘operative dimensions of capital and capitalism ... “hit the ground”’ (Mezzadra & Neilson, 2019, p. 2), where the ‘operative surface’ (Mezzadra & Neilson, 2019, p. 3), across which capitalist processes go to work, is ‘neither merely terrain nor land’ (Mezzadra & Neilson, 2019, p. 3) but an enmeshing of ‘spatial, social, legal, and political formations’ (Mezzadra & Neilson, 2019, p. 3). Thus, following Friedrich and Hoel’s (2023) formulation, a ‘media operation’ is a ‘technologically mediated action or procedure where symbolic and virtual resources are gathered to effect changes in the physical environment’ (Friedrich & Hoel, 2023, p. 53). With AI and ML models increasingly constituting the general foundations of capitalist processes and labour—what Matteo Pasquinelli refers to as the ‘automation of automation’ (Pasquinelli, 2023, p. 246)—media operations can now be understood to be driving capitalism itself, enmeshing these various formations across its operative surface.

Friedrich and Hoel consider how ‘[r]esearch on media operations often requires what we call an *empirical encounter* through which media scholars make themselves familiar with the media application under investigation, and just as importantly, with its context of use’ (Friedrich & Hoel, 2023, p. 52, authors’ emphasis). What such an ‘empirical encounter’ requires, however, is some kind of organizational device to structure it. Friedrich and Hoel term such a device ‘operative moments’ (Friedrich & Hoel, 2023, p. 60) which, for Noortje Marres (2020), might be understood as a computational ‘situation’ or for Karin Knorr Cetina (2009), a ‘synthetic situation’. Marres affirms the ‘constitutive role of computational settings, like electronic trading platforms, and digital media architectures, such as Skype, in the organisation of situations’ (Marres, 2020, p. 5).

Whether a ‘moment’ or a ‘situation’, what is important is how systems are *constructed* and *calibrated* for use. Indeed, that whilst Friedrich and Hoel also acknowledge the need to ‘pay due attention to the situations in which media operations unfold’ (Friedrich & Hoel, 2023, p. 57), Marres (2020) is more explicit in how situations are *actively*, and *knowingly*, cultivated. In this, operative moments do not magically or passively occur but are conscientiously generated with purpose and intent. In other words, that, as the book will explore, practitioners are actively involved in

developing strategies to generate specific kinds of (calibrated) situations, cultivating new operational conditions. This purposefulness is intrinsic to how decisions are made, and decision-making proceeds.

Adrian Mackenzie, in his detailing of the work of ‘machine learners’, also develops an approach he calls an ‘archaeology of operations’ (Mackenzie, 2017, p. 9), paying ‘attention to the specificity of practices’ that ‘is an elementary pre-requisite to understanding human-machine relations and their transformations’ (Mackenzie, 2017, p. 9). Attending to these operations allows Mackenzie to consider machine learning ‘as a form of knowledge production and a strategy of power’ (Mackenzie, 2017, p. 9) in the way that it carves out categories and differences through classificatory processes. A repeat interest throughout this book is the manner in which such cleavages are made in the domain of autonomous driving. These classificatory operations underpin the decision-making power of autonomous vehicles.

As Chap. 4 will detail, any operational analysis of autonomous driving must start from acknowledging its *interoperability*, and thus that any operational analysis necessarily demands an *interoperational* analysis. This matters because it shifts a focus away from the centre towards the edges, from systems or system components or modules to system connections: those critical contact points—especially so in automated and autonomous systems—that demand but don’t always offer stability, security, or seamlessness. It is at these critical contact points, as the book will explore, where decisions crystallize.

Methodologically, the decision can be understood as a specific computationally organized situation composed for the purposes of executing an action or set of actions. As an adaptation of Marres (2020), I am specifically interested, therefore, in how decision-making situations are enacted, and how these various mapping, training, sensing, planning, securing, relaxing, executing, and resisting phases are integral to the operation, and interoperability, of autonomous vehicles. The decision, therefore, is an entry point or cut into these operational phases, allowing one to bring these various components of data, model, operation, practice, path, network, and system together. Through the making of decisions, the actualization of the qualities of the autonomous vehicle becomes observable, at least provisionally.

In this, the book does not offer a straightforward ethnography of autonomous vehicle design or development but a methodological perspective centred on articulating how decision-making is practised, performed, and refined in a machine vision context. This ‘operational middle-range approach’ (Friedrich & Hoel, 2023, p. 52) is taken to best ameliorate the limitations of operational media, images, and data themselves, ordinarily preventing the kind of accessibility afforded by representational media.

‘Pipelines’

Typically, an operation might be understood as ‘a straightforward relation of cause and effect, input and output, that can be tracked according to the model of classical mechanics’ (Mezzadra & Neilson, 2019, p. 67). The idea of the *pipeline* is a critical organizational tool and metaphor for the development of autonomous vehicles, following the above, as witnessed in the flowchart of the software system underpinning the winning vehicle in the 2007 DARPA Urban Challenge, ‘Stanley’ (Fig. 1.2).

Thus, the ‘pipeline’ is used within the book as a methodological device in order to follow the sequential stages along which various decisions are carried out, from mapping, training, sensing, and simulating, to securing, relaxing, and resisting. However, as Thylstrup suggests, ‘traceability’ is a ‘mechanism that can allow individuals and organisations to forensically “follow” data along a life cycle or value chain invok[ing] an imaginary of linearity’ standing ‘at odds with the cultural complexities of the ethics of machine learning and the role of data sets within it’ (Thylstrup, 2022, p. 664). Whilst data might be understood to flow seamlessly from one stage to another, and that with this seamless flow the significance of each stage can simply be ‘read off’, such linearity is indeed a construction. As Mezzadra and Neilson agree:

However automatic or given the results of an operation might appear to be, when we look at it from this perspective, there is always a back story, a drama of fictions and tensions in which the efficacy of the operation

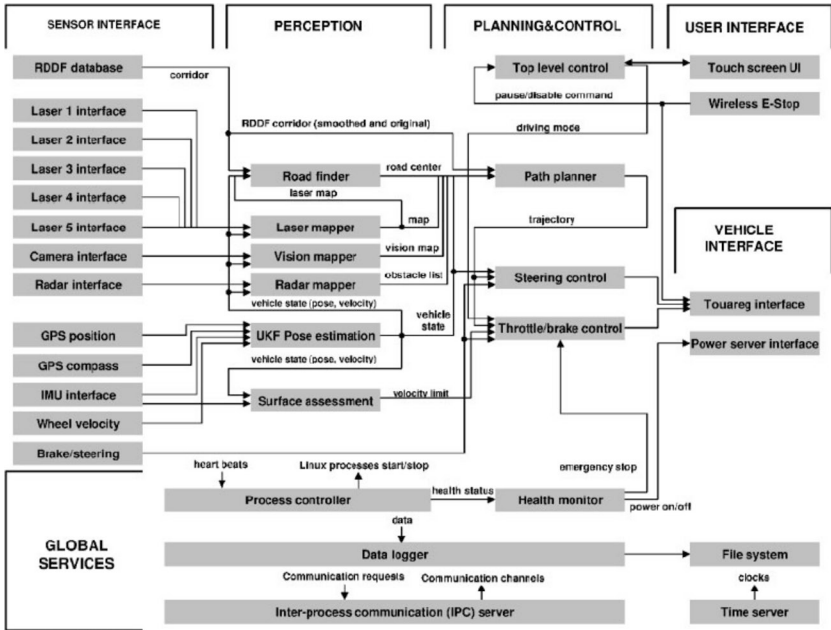


Fig. 1.2 Stanley processing 'pipeline'. (Source: Thrun et al., 2007)

appears far more fragile and elusive than might otherwise be assumed. (Mezzadra & Neilson, 2019, p. 67)

The chapters within attest to the methodological artifice of the pipeline metaphor, with some actors within the fields of machine learning and autonomous driving necessarily noting so. Neither in theory nor in practice is machine learning, or the automation of perception and decision-making, a linear process. As such, the book narrates the moments at which this linearity is deliberately challenged or otherwise threatened—the 'drama of fictions and tensions' Mezzadra and Neilson talk of. Indeed, that for some practitioners the enduring imaginary of the AI pipeline itself might, arguably, be the biggest threat to the delivery of autonomous vehicles altogether.

Phases

The book is not intended to be a ‘complete’ history of autonomous driving, nor really a history at all. Nevertheless, in the process, it has been conceptually valuable to develop an understanding of the eras or ‘phases’ of the development of autonomous driving. In this short history, the first DARPA Grand Challenge in 2004 is generally considered to have started the commercial race, funded by the US Defense Advanced Research Projects Agency (DARPA). In this early ‘robotic’ phase, autonomous driving was principally an off-road pursuit, centred around competing in these so-called grand challenges, and ensuring vehicles completed an assault course of sorts.

It is possible to trace the seeds of this initial phase of autonomous driving back even further to 1983, at the birth of a DARPA programme called the Strategic Computing Initiative (SCI), which ran for 10 years before being cancelled by the Clinton administration (Roland & Shiman, 2002). Designed to deliver ‘machine intelligence’, the SCI incorporated multiple interconnected, parallel, projects intended to produce concrete military systems concerning autonomous movement, battlefield management, and simulation (Roland & Shiman, 2002). One such project was the Autonomous Land Vehicle (ALV) programme, an early autonomous military vehicle, designed to ‘take the next step to high-level, real-time, three-dimensional IU [image understanding]’ (Roland & Shiman, 2002, p. 220). Built through a partnership between various public and private actors—Carnegie Mellon University (CMU), General Electric, Honeywell, and Columbia University—the ALV was given its debut in May 1985, ‘successfully navigating a 1,016-metre course in 1,060 seconds’ (Roland & Shiman, 2002, p. 228), better than anything that had come before it. Despite the early success, the project was cancelled in 1988 as the SCI faced scrutiny over standardization and compatibility issues.

As Chap. 2 discusses, the 1990s can be considered a dormant period for autonomous driving, punctuated only by CMU’s continued development of their Navlab vehicles—successors, of a kind, to the ALV (Roland

& Shiman, 2002). Only in the 2000s did interest return: the 2004 DARPA Grand Challenge—dubbed the first ‘great robot race’ (Buehler et al., 2007)—saw 106 original applicants each tasked with navigating a 150-mile off-road course in the Mojave Desert in California. In the end only 15 teams participated for the US\$1million cash prize, with CMU’s Red Team vehicle ‘Sandstorm’ winning, despite only managing to negotiate 5% of the total distance. Subsequently referred to as the ‘debacle in the desert’ (Hooper, 2004, p. n.p.), it was far from a success, either for DARPA or for the winning team, notably a previous partner in the ALV programme.

Despite this, two further iterations were organized, including a 2005 edition in which five teams managed to complete a 132-mile course across similar desert terrain, with Stanford University’s Stanley Racing Team winning in 6 hours 53 minutes with Stanley, and teams from CMU finishing in second and third (Buehler et al., 2007; Thrun et al., 2007). Following a two-year break, the competition re-emerged as the 2007 DARPA Urban Challenge, with teams now competing to complete an ‘urban’ rather than off-road course, including obeying all California state driving regulations. CMU’s Tartan Racing vehicle ‘Boss’ completed the course quickest, with Stanford’s Racing’s ‘Junior’ vehicle finishing second. This was the first time, as DARPA remarked, that ‘autonomous vehicles [had] interacted with both manned and unmanned vehicle traffic in an urban environment’ (DARPA, 2007, p. n.p.). Across three iterations, only teams from either CMU or Stanford had claimed the first prize.

The second ‘benchmark’ phase saw a shift away from robotics research towards the setting of technical ‘benchmarks’ to establish the commercial development of autonomous vehicles and ensure the comparability of autonomous driving methods. Key moments in this era that ran from 2009 to 2018 include the launch of Google’s self-driving car project in January 2009 (led by Chris Urmson, Team Tech Lead of Tartan Racing, winners of the 2007 DARPA Urban Challenge, assisted by Sebastian Thrun, Stanley Racing Team’s lead), the release of the Kitty Vision Benchmark Suite (2012), and the launch of Comma openpilot (2016) (Table 1.1).

The third, ongoing ‘incremental’ phase, is committed to the incremental development of ML methods used for autonomous driving. Key

moments in this period include the release of Waymo Open Dataset in August 2019, the launch of Waymo ‘Virtual Challenges’ in March 2020, and the acquisition of Uber ATG by Aurora (now lead by Chris Urmson) in December 2020. Cruise’s suspension of robotaxi services in San Francisco (2023) might be seen as an end to this incremental phase, as I detail in the final chapter.

The book also narrates the transition across a threshold in the development of autonomous vehicles: before the fatal Uber ATG crash in Tempe, Arizona, on March 18, 2018, that killed Elaine Herzberg, and after. It is this moment that I argue marks the end of the benchmark era and the beginning of the incremental era. Herzberg’s death was the first recorded fatality involving an autonomous vehicle, anywhere in the world. A full investigation was conducted by the US National Transportation Safety Board (NTSB) into the cause of the crash, producing a final report on

Table 1.1 Key dates in the development of autonomous driving, 1985–2023

Date	Event	Era	Key term	Notes
May 1985	Autonomous Land Vehicle (ALV) demonstration	Robotic	Strategic computing	ALV project cancelled in 1988
March 13, 2004	First DARPA Grand Challenge		Terrain	No winners
October 8, 2005	Second DARPA Grand Challenge			Won by Stanford Racing Team, led by Sebastian Thrun, ML-driven approach
November 3, 2007	DARPA Urban Challenge		Urban	Won by Tartan Racing, Chris Urmson Team Tech Lead, funded by CMU/GM/AnnieWAY by KIT DNF

(continued)

Table 1.1 (continued)

Date	Event	Era	Key term	Notes
January 17, 2009	Google self-driving car project launches	Benchmark	Project	Led by Urmson
March 20, 2012	Kitti Vision Benchmark Suite released		Street benchmarking	Co-author was Raquel Urtasun, co-developed by KIT team
February 2, 2015	Uber and Carnegie Mellon University announce 'partnership'		Strategic collaboration	Uber acquire talent <i>en masse</i> from CMU Robotics Institute
February 20, 2016	Cityscapes dataset released		Semantic	TU Darmstadt, Max Planck Institute for Informatics, Daimler AG collaboration
September 13, 2016	Comma One launched at Disrupt SF		Open	Cancelled a month later, after NHTSA special order
November 30, 2016	Comma openpilot launched			Publicly released on Github
December 13, 2016	Waymo launches		Platform	Official spin-off from in-house Google project
March 19, 2018	Fatal Uber ATG crash in Tempe, Arizona		Street testing	Raquel Urtasun was Uber ATG Chief Scientist and Head of Research

(continued)

Table 1.1 (continued)

Date	Event	Era	Key term	Notes
August 21, 2019	Waymo Open Dataset released	Incremental	Data	First release of training dataset by solo commercial operator
March 19, 2020	Waymo OD Virtual Challenges launched		Competition	First public challenges by solo commercial operator
December 7, 2020	Aurora acquires Uber ATG		Consolidation	Aurora founded by Urmson
June 14, 2021	Waabi launched		Simulation	Founded by Urtasun
July 31, 2021	Comma three launched		Device class	Compatible with over 200 vehicle models
June 28, 2022	Major Cruise outages in San Francisco		Live	Greatest public fleet-wide technical error to date
August 10, 2023	CPUC approves Cruise operations without restriction in SF		Authorization	CPUC regulates services and utilities across California
October 24, 2023	California DMV suspends Cruise operation permits in SF		Suspension	California DMV registers motor vehicles and issues driving licences across California

November 19, 2019 (NTSB, [2019](#)). On September 15, 2020, the vehicle operator (VO) behind the wheel of the vehicle at the time of the crash, Rafaela Vasquez, was charged with negligent homicide (Dungan, [2020](#)). On December 7, 2020, Uber agreed a deal to sell their autonomous vehicle operations, Uber ATG, to Aurora—a fellow autonomous vehicle firm—valuing it at \$10billion (Korosec, [2020](#)). The book thus argues in different ways that the development of autonomous vehicles has been

fundamentally changed by the events in Tempe in 2018, with continuing operators forced into various levels of introspection. Hence, the rise in incrementalist thinking within the industry after 2018, and the birth of a new era of autonomous driving.

Nonetheless, before Tempe and after Tempe are not especially sharp distinctions. Key events following the crash itself have reshaped the field along the way, from the publication of a preliminary report into the crash by the NTSB on May 24, 2018 (Krishna, 2018) followed by the final report 18 months later, then the decision to charge Vasquez 10 months after, arguably culminating in the sell-off of Uber ATG at the end of 2020 (Smiley, 2022, 2023). Connectedly, the crash, reports, indictment, and sale constitute significant moments at which the trajectory—and ultimately, the promise—of autonomous driving changed course. The book is thus an attempt to narrate this course, as well as exploring the other parallel trajectories along which fellow operators have themselves travelled. Chapters 5, 6, 7, and 8 all concern the consequences of the fatal crash in various ways, focusing on technical developments, key actors, and public discourses that I argue would not have occurred without the crash on that fateful day in Tempe, Arizona, in 2018.

Chapter Outline

As the start of this introduction contended, the book offers an insight into the ‘advance decisions’ deemed necessary to automate driving. In this, it is worth establishing what the book *isn’t* about. Firstly, it is not really about the automotive industry. From early in the research process, it became apparent that the quest to automate driving was not being conducted by traditional automotive manufacturers. Wedded to established relationships with technology suppliers and bound by safety regulations and common industry practices, car manufacturers were neither interested in offering, nor able to offer, the dream of autonomous driving. Chapter 6, on the semiconductor ‘chip crisis’, is the only exception.

Whilst many advanced driver-assistance systems (ADAS) have been launched by car manufacturers over recent years, including Volvo’s vaunted ‘moose vision’ (Adams, 2017), these incremental systems have

not been the landmark, game-changing, epoch-defining, advancements promised by a fully autonomous future. Instead, they were decidedly incremental offers: the kinds of systems used to tempt customers to buy upgrades to base vehicle models on display. Despite genuinely intriguing speculative proposals, such as Nissan's Seamless Autonomous Mobility project (Hind, 2022), most automotive manufacturers stuck to what they knew: a combination of aforementioned ADAS alongside concept cars. Rarely ever constituting concrete blueprint for future vehicle models, concept cars are used by manufacturers as embodiments of future-forward thinking—the direction of travel for the automotive brand, and little else. This mobilization of a 'technological sublime' (Hildebrand, 2019) is part of the standard automotive playbook.

Secondly, the book isn't really about autonomous vehicles, per se—and certainly not about 'cars'. No single chapter examines any actually existing autonomous vehicle, what it is, and what it can do. More accurately, the book is about the myriad forms of technical work being executed to ideally, possibly, hopefully deliver a form of autonomous driving in the future. Even the final chapter on autonomous vehicle passenger services (AVPS) in San Francisco is not really concerned with the vehicle as a singular, material object—but the regulatory and practical battles concerning it. Even more so, the book does not intend to tell a story about cars in any meaningful sense. Despite a sustained theoretical focus on decision-making—all leading up to the moment of executing live manoeuvres on the road itself—the book deliberately focuses on the different kinds of (technical) decisions being made in advance of such moments. Indeed, rather than just a heuristic device, the book mirrors the material efforts of engineers, annotators, vehicle operators, regulators, enthusiasts, and activists on the ground—few of which (besides avid participants in the Comma project or test vehicle operators) are actually ever concerned with, or get close to, the act of driving itself. Even in such cases, the car itself recedes from view, dissolved into material components and systems each tweakable, modifiable, and optimizable to offer a greater affective 'autonomous' driving experience.

Chapter 2 considers the role of mapping in the development of autonomous vehicles, offering a comparison between two firms—Uber ATG and Waabi—taking different paths to establishing the cartographic

coordinates of their respective enterprises. Considered a fundament for such work, autonomous vehicle developers must concern themselves with building ‘operational design domains’ (ODDs), mapping the full geographic extent of their autonomous vehicle operations, down to the final millimetre. Far from a purely scientific endeavour, defined only by safety protocols or technical requirements, the setting of ODDs—most certainly in the case of Uber ATG—are informed by economic decisions concerning viability, scalability, and profitability. This fastidious ODD work—manual route driving, data tagging, domain building—crystallizes the search for viable future autonomous ride-hailing markets. Other operators, however, offer a different cartographic approach, less concerned with the mapping of geometric phenomena. Waabi’s stated ‘map-less’ approach to autonomous driving, although far from a cartographic, demonstrates a sensor data-driven approach seeking to eradicate pesky errors that enter into the developmental pipeline at the mapping stage. These attempts to offer innovative solutions to traditional engineering problems—such as the modular, ‘software stack’ approach—are typical of the competitive nature of autonomous driving.

Chapter 3 examines two training datasets—the KITTI Vision Benchmark Suite and Waymo Open Dataset—constituting key milestones in the ‘ground truthing’ of autonomous vehicles. Launched only seven years apart, each dataset is nonetheless markedly different, constitutive of very different moments in the development of autonomous vehicles—what I have called here the benchmark and incremental phases of autonomous driving. Often taken-for-granted, training datasets are integral components in the development of autonomous vehicles, critical to the building of ML models that underpin their executive functions. Only ever as useful as the classification and annotation work performed on them, autonomous vehicle training datasets—such as the KITTI Vision Benchmark Suite and Waymo Open Dataset—demonstrate great variety in composition, technical set-up, source data, and volume, amongst many other things. The KITTI Vision Benchmark Suite, launched in 2012, offered the first real-world benchmarks for the comparison of machine vision systems used in autonomous driving settings. An emerging interest in ‘interesting-ness’, typified by Waymo engineers’ repeated emphasis on scenario diversity within such datasets, similarly

defines the incremental era of autonomous driving, committed to refining extant machine learning and machine vision techniques.

Chapter 4 explores the ‘sensor strategies’ being devised to optimize object-recognition processes, critical to gifting autonomous vehicles their perceptive qualities. Although such sensor work is already necessary for the capture of ML training data in the first instance, novel methods and approaches are likewise required in order to improve what researchers call the ‘online’ (i.e. real-time) capabilities of their object-recognition systems. Drawing on work on the ‘operational’ nature of digital images and data, the chapter considers how machine vision researchers, such as those at Argo AI, are engaged in various efforts to rationalize the object-recognition process. Offering practical examples of the need for ‘interoperability’ between different stages in the autonomous vehicle decision-making process, such work demonstrates the need for interoperability on different technical, epistemological, and organizational levels. Practical examples encountered during autonomous driving and machine vision workshops include innovative, interstitial techniques to upgrade lidar to a fully 3D sensing format, and ‘dynamic scheduling’ techniques to balance quick and accurate image understanding. In both cases the chapter understands such work as integral to ‘finessing’ the interoperability of autonomous vehicle systems, ensuring ingested sensor data is prepared for subsequent stages in the decision-making pipeline.

Chapter 5 details the work undertaken by Waymo—Google/Alphabet’s autonomous vehicle division—to demonstrate the safety of their autonomous vehicles. Through analysis of two documents, the chapter considers how Waymo has publicly reported so-called contact events—the company’s euphemism for incidents and crashes involving Waymo vehicles. Beginning in 2020, Waymo published a report detailing 47 contact events involving Waymo vehicles undergoing testing on public roads. Twenty-nine were considered ‘simulated events’, incidents that would have happened if a human operator had not assumed control. Generated instead within in-house simulation software, Waymo’s revealing of such work constituted an important milestone in the post-Uber ATG crash autonomous world, suggesting that the incident had forced a change in Waymo’s public documentation of the crash worthiness of their own vehicles. A subsequent report in 2023 further demonstrated Waymo’s

efforts to document their own safety protocols around the testing of autonomous vehicles. Centred only on actual contact events (rather than simulated ones), the report chronicles 20 incidents that Waymo vehicles were implicated in from September 2020 to January 2023. Derived from two different test locations—Phoenix, Arizona, and San Francisco—the data is intended to show the lengths the company has gone to classify contact events, developing a ‘conflict typology’ comprising of 16 such types of possible interactions between the Waymo vehicle and other road users. The chapter argues that Waymo, through such classification work, is effectively able to apportion blame away from their own vehicle, surmising that all other road users (and sometimes, inanimate objects) are responsible for the recorded encounters. As examples of ‘deferred’ decision-making—shifting the meaning of such incidents into the future, and away from Waymo—algorithmic ‘doubt’ is used as an eminently useful, reputational resource to control the significance of such events as documented, discussed, and circulated in the public arena.

Chapter 6 explores how different critical components of autonomous vehicles are ‘secured’ through a range of techniques. In the first instance, the chapter considers how the decision-making capacities of autonomous vehicles are dependent on them operating as ‘unfamiliar’ forms of sovereign actors, following Bratton (2016). In such cases, the autonomous vehicle offers a kind of ‘functional sovereignty’ (Pasquale, 2017), derived through the software that underpin them, to which other parties are subject. Firms in such cases, as becomes apparent in subsequent chapters, have been able to carefully manage access to underlying decision-making data—even if compelled to release it for regulatory reasons. The second part of the chapter focuses on efforts to secure a piece of computational hardware critical for the development of autonomous vehicles: the semiconductor chip. The ongoing ‘chipification’ (Forelle, 2022) of automobility has led to an increased demand for semiconductors within the automotive sector over the past few years, culminating in a Covid-19 pandemic induced ‘chip crisis’. Exposing vulnerabilities in the global supply chain of the semiconductors, the chip crisis surfaced reliance on one particular company: the Taiwan Semiconductor Manufacturing Company (TSMC), the world’s leading chip fabricator. Through an analysis of TSMC and the global semiconductor chip market, the chapter

considers how the European Union (EU) has relaxed rules around state aid intended to boost future capacity in chip fabrication in the EU bloc. Dependent upon a form of economic exceptionalism unique to eastern Germany, the chapter examines how firms such as Intel are ‘onshoring’ chip production for the benefit of (German) carmakers.

Chapter 7 shows how an open-source driver-assistance start-up called Comma is attempting to ‘make driving chill’ through a form of empowered individualism. Understood as a contemporary manifestation of the ‘Californian ideology’ (Barbrook & Cameron, 1996), the Comma project offers a rather unique alternative to either a tech-driven dream of autonomous driving or the automotive industries’ offer of assistive driving. Fostering an active, online community of users, Comma can be understood as an interesting mutation of open-source ethics, a derivation of McKenzie Wark’s (2004) idea of the ‘hacker class’. Cultivating a wide range of antagonists—from ‘conehead’ engineers at Waymo, to traditional car manufacturers—the Comma community fosters a kind of ‘spiritual communion’, in which members share experiences of using their Comma devices. Committed to a ‘culture of testing’, members happily compose notes on how to calibrate their devices, report unknown device bugs, document their experiences in video and photographic form, and ultimately, take pleasure from being in control of the automation of their own vehicle. In such instances, as the chapter contends, Comma draws on, and extends, an electronic form of libertarianism lost to platform capitalism.

Chapter 8 considers the nascent rise of anti-autonomous vehicle activism in San Francisco. Centred on Cruise’s lengthy battle to receive permission to operate autonomous vehicle passenger services (AVPS), the chapter considers how resistance to autonomous vehicles has slowly fermented. Connecting opposition back to the infamous protests against ‘Google buses’ shuttling workers to big tech campuses in the early 2010s, the chapter narrates the story of initial regulatory forms of resistance to AVPS operations in San Francisco. After calls by municipal transportation bodies for an ‘incremental’ approach to the public rollout of AVPS operations, resistance abruptly moved out of a regulatory domain in July 2023 with the launch of a campaign by an activist group called Safe Street Rebel. During a ‘Week of Cone’ called for by the group, Cruise

vehicles became paralysed by humble traffic cones placed on the hood of each vehicle. Despite such efforts—the first, sustained street-based opposition to autonomous vehicles anywhere in the world—after a six-hour hearing, Cruise’s application for expanded AVPS operations were approved—only for them to be swiftly suspended following implication in a road traffic accident. The following month, Cruise CEO Kyle Vogt resigned, bringing an apparent end to a tumultuous period in the rollout of AVPS operations in San Francisco—and possibly, to autonomous vehicles, full stop.

Together, the book offers a composite insight into the multiplicity of efforts to deliver autonomous driving—in some form, somewhere—to the general public. In so doing, it offers an opportunity to consider how these efforts have required huge amounts of investment and labour, without much to show for it—with numerous actual and reputational casualties along the way. Affording a critical view of the last 10 years of the desire to automate driving offers an insight into the efforts big tech firms and comparative start-ups have gone to apply AI and ML techniques to the human, thoroughly modern, and undeniably routine task of driving.

References

- Adams, E. (2017, January 27). Volvo’s cars now spot moose and hit the brakes for you. *Wired*.
- Alpaydin, E. (2016). *Machine learning: The new AI*. MIT Press.
- Amoore, L. (2020). *Cloud ethics: Algorithms and the attributes of ourselves and others*. Duke University Press. <https://doi.org/10.2307/j.ctv11g97wm>
- Amoore, L. (2022). Machine learning political orders. *Review of International Studies*, 49(1), 20–36. <https://doi.org/10.1017/S0260210522000031>
- Andrejevic, M. (2019). *Automated media*. Routledge.
- Bader, V., & Kaiser, S. (2019). Algorithmic decision-making? The user interface and its role for human involvement in decisions supported by artificial intelligence. *Organization*, 26(5), 655–672. <https://doi.org/10.1177/1350508419855714>
- Bratton, B. (2016). *The stack: On software and sovereignty*. MIT Press.
- Beer, S. (1959). *Cybernetics and management*. The English Universities Press.
- Beer, S. (1981 [1972]). *Brain of the firm*. John Wiley & Sons.

- Beer, S. (1995 [1966]). *Decision and control: The meaning of operational research and management cybernetics*. Wiley.
- Bessner, D., & Guilhot, N. (Eds.). (2019). *The decisionist imagination: Sovereignty, social science and democracy in the 20th century*. Berghahn Books. <https://doi.org/10.2307/j.ctvw04b7q>
- Bucher, T. (2018). *If... Then: Algorithmic power and politics*. Oxford University Press. <https://doi.org/10.1093/oso/9780190493028.001.0001>
- Buehler, M., Iagnemma, K., & Singh, S. (Eds.). (2007). *The 2005 DARPA grand challenge: The great robot race*. Springer.
- DARPA. (2007). *DARPA urban challenge*. DARPA. <https://www.darpa.mil/about-us/timeline/darpa-urban-challenge>
- Dencik, L., & Stevens, S. (2021). Regimes of justification in the datafied workplace: The case of hiring. *New Media & Society*, 25(12), 3657–3675. <https://doi.org/10.1177/14614448211052893>
- Derrida, J. (1992). Force of law: ‘The mystical foundation of authority’. In D. Cornell, M. Rosenfeld, & D. Carlson (Eds.), *Deconstruction and the possibility of justice* (pp. 3–67). Routledge.
- Dungan, R. (2020, September 15). *Negligent homicide charges filed in Tempe self-driving Uber death*. KJZZ. <https://kjzz.org/content/1618588/negligent-homicide-charges-filed-tempe-self-driving-uber-death>
- Dwyer, A. (2020, May 8). *Algorithms don't make decisions!* Digital societies. <https://www.digitalsocieties.co.uk/blog/algorithms-dont-make-desicions>
- Farocki, H. (2004). Phantom images. *Public*, 29, 12–22.
- Friedrich, K., & Hoel, A. A. S. (2023). Operational analysis: A method for observing and analyzing digital media operations. *New Media & Society*, 25(1), 50–71. <https://doi.org/10.1177/1461444821998645>
- Galloway, A. (2004). *Protocol: How control exists after decentralization*. MIT Press.
- Hayles, N. K. (2017). *Unthought: The power of the nonconscious*. The University of Chicago Press.
- Hildebrand, J. M. (2019). On self-driving cars as a technological sublime. *Techné: Research in Philosophy and Technology*, 23(2), 153–173. <https://doi.org/10.5840/techné2019722100>
- Hind, S. (2022). Making decisions: The normal interventions of Nissan ‘mobility managers’. *Mobilities*, 17(4), 467–483. <https://doi.org/10.1080/17450101.2021.1988682>
- Hind, S., van der Vlist, F., & Kanderske, M. (2024). Challenges as catalysts: How Waymo’s Open Dataset Challenges shape AI development. *AI & Society*, 0(0), 1–17. <https://doi.org/10.1007/s00146-024-01927-x>

- Hong, S.-h. (2022b). Predictions without futures. *History and Theory*, 61(3), 371–390. <https://doi.org/10.1111/hith.12269>
- Hong, S.-h. (2022a). Prediction as extraction of discretion. *2022 ACM Conference on Fairness, Accountability, and Transparency (FAccT '22)*, June 21–24, 2022, Seoul, Republic of Korea ACM, New York, USA: 1–10. <https://doi.org/10.1145/3531146.3533155>
- Hooper, J. (2004, June 4). From DARPA grand challenge 2004: DARPA's debacle in the desert. *Popular Science*. <https://www.popsoci.com/scitech/article/2004-06/darpa-grand-challenge-2004darpas-debacle-desert/>
- Jaton, F. (2021). *The constitution of algorithms: Ground-truthing, programming, formulating*. MIT Press. <https://doi.org/10.7551/mitpress/12517.001.0001>
- Kien, G. (2008). Technography = technology + ethnography: An introduction. *Qualitative Inquiry*, 14(7), 1101–1109. <https://doi.org/10.1177/1077800408318433>
- Knorr Cetina, K. (2009). The synthetic situation: Interactionism for a global world. *Symbolic Interaction*, 32(1), 61–87. <https://doi.org/10.1525/si.2009.32.1.61>
- Korosec, K. (2020, December 7). Uber sells self-driving unit Uber ATG in deal that will push Aurora's valuation to \$10B. *Techcrunch*. <https://techcrunch.com/2020/12/07/uber-sells-self-driving-unit-uber-atg-in-deal-that-will-push-auroras-valuation-to-10b/>
- Krishna, S. (2018, May 24). *NTSB's preliminary report on Uber crash focuses on emergency braking*. Engadget. <https://www.engadget.com/2018-05-24-ntsb-preliminary-report-uber-self-driving-crash-emergency-braking.html>
- Latour, B. (1986). Visualisation and cognition: Drawing things together. In H. Kuklick (Ed.), *Knowledge and society studies in the sociology of cultures past and present* (pp. 1–40). Jai Press.
- Law, J. (2002). *Aircraft stories: Decentering the object in technoscience*. Duke University Press.
- Leigh Star, S. (2010). This is not a boundary object: Reflections on the origin of a concept. *Science, Technology, & Human Values*, 35(5), 601–617. <https://doi.org/10.1177/0162243910377624>
- Leigh Star, S., & Griesemer, J. R. (1989). Institutional ecology, 'translations' and boundary objects: Amateurs and professionals in Berkeley's Museum of Vertebrate Zoology, 1907–39. *Social Studies of Science*, 19(3), 387–420. <https://www.jstor.org/stable/285080>
- Mackenzie, A. (2017). *Machine learners: Archaeology of a data practice*. MIT Press.

- Marres, N. (2020). For a situational analytics: An interpretive methodology for the study of situations in computational settings. *Big Data & Society*, 7(2), 1–16. <https://doi.org/10.1177/2053951720949571>
- Mendon-Plasek, A. (2020). Mechanized significance and machine learning: Why it became thinkable and preferable to teach machines to judge the world. In J. Roberge & M. Castelle (Eds.), *The cultural life of machine learning: An incursion into critical AI studies* (pp. 31–78). Palgrave Macmillan. https://doi.org/10.1007/978-3-030-56286-1_2
- Mezzadra, S., & Neilson, B. (2019). *The politics of operations: Excavating contemporary capitalism*. Duke University Press.
- National Transportation Safety Board (NTSB). (2019, May 11). *Collision between vehicle controlled by developmental automated driving system and pedestrian, Tempe, Arizona, March 18, 2018*. NTSB. Retrieved October 1, 2023 from <https://data.nts.gov/Docket/Document/docBLOB?ID=40479021&FileExtension=.PDF&FileName=NTSB%20-%20Adopted%20Board%20Report%20HAR-19%2F03-Master.PDF>
- Parisi, L. (2017, October 1). Reprogramming decisionism. *E-flux*. <https://www.e-flux.com/journal/85/155472/reprogramming-decisionism/>
- Pasquale, F. (2017, December 6). From territorial to functional sovereignty. *Law and Political Economy*. <https://lpeproject.org/blog/from-territorial-to-functional-sovereignty-the-case-of-amazon/>
- Pasquinelli, M. (2021). How to make a class: Hayek's neoliberalism and the origins of connectionism. *Qui Parle*, 30(1), 159–184. <https://doi.org/10.1215/10418385-8955836>
- Pasquinelli, M. (2023). *The eye of the master: A social history of artificial intelligence*. Verso.
- Rella, L. (2023). Close to the metal: Towards a material political economy of the epistemology of computation. *Social Studies of Science*, 0(0), 1–27. <https://doi.org/10.1177/03063127231185095>
- Ribes, D., Hoffman, A. S., Slota, S. C., & Bowker, G. C. (2019). The logic of domains. *Social Studies of Science*, 49(3), 281–309. <https://doi.org/10.1177/0306312719849709>
- Rieder, B. (2020). *Engines of order: A mechanology of algorithmic techniques*. Amsterdam University Press. <https://doi.org/10.2307/j.ctv12sdvfl>
- Rieder, B., & Skop, Y. (2021). The fabrics of machine moderation: Studying the technical, normative, and organizational structure of Perspective API. *Big Data & Society*, 8(2), 1–16. <https://doi.org/10.1177/205395172111046181>

- Roberge, J., & Castelle, M. (2020). Toward an end-to-end sociology of 21st-century machine learning. In J. Roberge & M. Castelle (Eds.), *The cultural life of machine learning: An incursion into critical AI studies* (pp. 1–30). Palgrave Macmillan. https://doi.org/10.1007/978-3-030-56286-1_1
- Roland, A., & Shiman, P. (2002). *Strategic computing: DARPA and the quest for machine intelligence*. MIT Press.
- Sigaut, F. (1993). Technology. In T. Ingold (Ed.), *Companion encyclopedia of anthropology: Humanity, culture and social life* (pp. 420–459). Routledge.
- Simondon, G. (2017). *On the mode of existence of technical objects*. Univocal Publishing.
- Smiley, L. (2022, March 8). ‘I’m the operator’: The aftermath of a self-driving tragedy. *Wired*. <https://www.wired.com/story/uber-self-driving-car-fatal-crash/>
- Smiley, L. (2023, July 28). The legal saga of Uber’s fatal self-driving car crash is over. *Wired*. <https://www.wired.com/story/ubers-fatal-self-driving-car-crash-saga-over-operator-avoids-prison/>
- Sprenger, F. (2015). *The politics of micro-decisions: Edward Snowden, Net neutrality, and the architectures of the Internet*. Meson Press. <https://doi.org/10.14619/005>
- Sprenger, F. (2020). Microdecisions and autonomy in self-driving cars: Virtual probabilities. *AI & Society*, 37, 619–634. <https://doi.org/10.1007/s00146-020-01115-7>
- Sprenger, F. (2021). Intervals of intervention: Micro-decisions and the temporal autonomy of self-driving cars. In A. Volmar & K. Stine (Eds.), *Media infrastructures and the politics of digital time: Essays on hardwired temporalities* (pp. 157–174). Amsterdam University Press. <https://doi.org/10.2307/j.ctv1xcxr3n.12>
- Steinhoff, J., & Hind, S. (2024). Simulation and the reality gap: Moments in a prehistory of synthetic data. *MediArXiv*. <https://doi.org/10.33767/osf.io/np3vb>
- Theureau, J. (2003). Course-of-action analysis and course-of-action centered design. In E. Hollnagel (Ed.), *Handbook of cognitive task design* (pp. 55–81). Lawrence Erlbaum.
- Thrun, S., Montemerlo, M., Dahlkamp, H., Stavens, D., Aron, A., Diebel, J., et al. (2007). Stanley: The robot that won the DARPA Grand Challenge. In M. Buehler, K. Iagnemma, & S. Singh (Eds.), *The 2005 DARPA grand challenge: The great robot race* (pp. 1–43). Springer.

- Thylstrup, N. B. (2022). The ethics and politics of data sets in the age of machine learning: Deleting traces and encountering remains. *Media, Culture & Society*, 44(4), 655–671. <https://doi.org/10.1177/01634437211060226>
- Van der Vlist, F., Helmond, A., & Ferrari, F. (2024). Big AI: Cloud infrastructure dependence and the industrialisation of artificial intelligence. *Big Data & Society*, 11(1), 1–16. <https://doi.org/10.1177/20539517241232630>
- Walker Rettburg, J. (2020). Situated data analysis: A new method for analysing encoded power relationships in social media platforms and apps. *Humanities & Social Sciences Communications*, 7(5), 1–13. <https://doi.org/10.1057/s41599-020-0495-3>
- Wark, M. (2004). *The hacker manifesto*. Harvard University Press.
- Wiener, N. (1973 [1948]). *Cybernetics*. MIT Press.
- Woolgar, S. (1998). A new theory of innovation? *Prometheus: Critical Studies in Innovation*, 16(4), 441–452. <https://doi.org/10.1080/08109029808629293>



2

Mapping Decisions: The Promise of Mapless-ness

This chapter examines how mapping decisions are made and the changing nature of mapmaking and cartographic data collection in the world of autonomous driving. Whilst the automation of driving also likely spells the end of automotive navigation in its current form, the chapter considers how curious innovations in the production, and collection, of cartographic data are critical to delivering autonomous vehicles.

The chapter distinguishes between two types of cartographic practice: the setting of operational design domains (ODDs) and the collection of semantic user data. On the one hand, geometric data offers a granular record of fixed objects within the driving environment, from lane widths to curb edges. On the other hand, semantic data provides an aggregate record of driving behaviour, from collated vehicle speeds to average stopping positions. Whilst the former is required to establish the outer (legal) limits for driving—the ODD—the latter is deemed necessary to establish, and ultimately emulate, human driving norms and behaviours. The chapter will focus on two cases: now-defunct Uber ATG's map-based, rule-dependent approach to autonomous driving and Waabi, an autonomous vehicle start-up led by former Uber ATG chief scientist, Raquel Urtasun who have developed the concept of 'mapless' driving (Casas et al., [2021](#)).

What is interesting about these cases is the extent to which the mapping of physical phenomena either containing specific road information (road signs, speed limit signs) or not (curbs, road markings) is far from the full extent of the mapping exploits undertaken by autonomous vehicle firms. Indeed, in both, a range of strategies to map *actions*, *movements*, *activities*, and *experiences* are witnessed—offering what fellow operator Mobileye calls a ‘semantic level’ understanding (Hind & Gekker, 2024). Each of these cases, representative of efforts more broadly, is therefore primarily interested in mapping the *social life* of the road. As a result, they must grapple with the social dynamics and particularities of driving, and the differences between how and where people actually drive, in order to provide a ‘ground truth’ for their autonomous vehicle systems. It is the mapping of these social, semantic elements that can be considered the firmer ‘ground’ for such truth-making practices—even more so than the fixed, geometric elements delineating the asphalt and concrete of the roads themselves.

This marks a shift in what cartography is and what kinds of cartographic work are being performed in the twenty-first century. In these cases, ‘mapless’ mapping is not only alive but has established itself—certainly within the domain of autonomous driving—as the principle form of cartographic work in the present. Mapless mapping, as the oxymoron suggests, isn’t interested in producing maps as readable outputs—or anything that especially looks or functions like one. Instead, mapless mapping is most interested in forming cartographic *impressions* based on the mapping of actions, movements, activities, and experiences. These cartographic impressions are designed to impress themselves *on*, and express themselves *in*, the machine learning (ML) models iteratively informing the future decision-making qualities of autonomous vehicles. Without doing so, they would fail to emulate the behaviours of the human drivers they wish to imitate.

Automotive Navigation

Whilst the humble road map, or a satnav, might be considered as singularly important to automotive navigation, the map itself has not always been integral to automotive navigation.

Tristan Thielmann, for instance, has considered the role of ‘photo-auto guides’ as early forms of ‘augmented’ automotive navigation devices (Thielmann, 2016). By superimposing orientational text and arrows on photographs or moving images, photo-auto guides (such as those produced by Rand McNally & Company) were intended to reassure drivers as to their movement along a designated, predetermined route. Designed to be used ‘en route’ rather than as reference maps, they can be considered non-digital ‘turn-taking machine[s]’ (Singh et al., 2019, p. 287) meant to alleviate ‘wayfinding troubles’ (Singh et al., 2019, p. 288) that might ordinarily arise during an unfamiliar journey.

As practical devices, only reaching their full utility upon a particular driving situation or route, photo-auto guides can be understood as ‘ontogenetic’, ‘enacted to solve relational problems’ (Kitchin & Dodge, 2007, p. 335) such as the “gap” between a plan for a route composed in advance and an ongoing journey that attempts to follow the plan’ (Singh et al., 2019, p. 288). A term favoured by French philosopher Gilbert Simondon (2017), ontogenesis refers to an object’s (contingent, speculative) unfolding or becoming, in contrast to the idea of ontology (being, form). In cases of navigational use, therefore, devices like the photo-auto guide can be thought of as enacting a generative possibility because, until the point of navigational requirement—and even within it—the guide itself has the capacity to actualize its navigational capacities or tendencies in particular ways, showing itself to be useful on this bend in the road or at that junction. If the guide contains a mistake or an imprecise description of the route ahead, then the guide’s navigational actualities might shrink or reform, or its capacities might fail to be realized at all.

As Thielmann suggests, these guides can be understood as ‘operative’ in two senses: firstly, in that the photographs themselves are ‘subjected to operative changes through information being embedded in their surface’, and secondly, in that the photographs are integrated into ‘part of an

operative practice: the practice of autonavigation' (Thielmann, 2016, p. 162). This operative practice of auto(motive) navigation is an example of the ontogenesis of cartographic calculation (Hind, 2020). In this, navigational practice is determined both by the 'capture' of objects or activities (subsequently generating geometric or semantic data), as well as the 'addition' of more, new, or different renderings of the world into it. In short, that any and every navigational device is constantly involved in the push-and-pull of extracting things out of, as well as adding things into, the world. Differences between navigational devices (paper maps, photo-auto guides, satnavs, map apps) therefore do not lie, *per se*, in whether one is nominally extractive or additive, but in how (and often where) these extractive and additive moments are executed.

Thus, whilst some might mourn the 'lost art' (Fisher, 2013, p. n.p.) of map reading with the rise of autonomous driving, the history of automotive navigation suggests that map reading itself has always shifted and morphed, both in terms of the constituent 'map' being read, and the map 'reading' as a specific set of skills, enacted whilst driving. Navigational information is still being captured, but in different ways, for different purposes and for different users. As such, automotive navigation can be divided into different temporal phases I call here: route, general, positional, and executive. These later eras (positional, executive) overlap to some degree with the phases of autonomous driving presented in the introduction (robotic, benchmark, incremental).

The early 'route' phase, from 1880 to 1945, can be considered the era of photo-auto guides and graphical road maps tied to specific routes or journeys from A to B (Schulten, 2000; Thielmann, 2016). The mid-phase from 1945 to 1990 can be considered the 'general' road atlas era, where entire national and international road atlases, complete with exhaustive visual and indexical records of roads and road-related services, could be found (Harley, 1989; Akerman, 2006; Wood, 2010). Both phases pre-date the development of autonomous vehicles in any real sense, although the emblematic navigational technologies of these eras—the photo-auto guide, the road atlas—can be considered extensions of the 'driver-car' (Dant, 2004) or 'car-driver' (Sheller & Urry, 2000) assemblage (Fig. 2.1).

The late 'positional' phase, from 1990 to 2009, can be considered the era of the satnav, as standalone and in-built satellite navigation systems

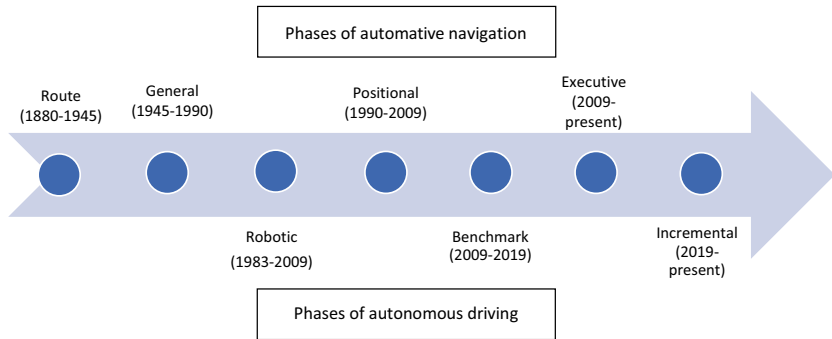


Fig. 2.1 A timeline of automotive navigation and autonomous driving

began to be integrated into cars (French, 2006; Wilken & Thomas, 2019; Hind, 2019). As the table in the introduction suggests, efforts to automate driving were largely dormant through the 1990s, whilst the satnav era of automotive navigation blossomed. The reasons were arguably twofold. Firstly, due to the more successful development of networking technologies. As Roland and Shiman wrote, early internet pioneers such as the US Defence Advanced Research Project Agency's (DARPA) Barry Leiner 'chose not to seek SC [Strategic Computing] money for networking' even though it was 'clearly infrastructure of the kind envisioned in the SC plan' (Roland and Shiman 2002, p. 110). DARPA's Strategic Computing Initiative (SCI), in which the Autonomous Land Vehicle (ALV) project was housed, did not involve the development of networking technologies that 'would ultimately contribute to the creation of the Internet' (Roland and Shiman 2002, p. 110). As the SCI fell apart at the end of the 1980s, interest in networking technologies soared, with development of the internet absorbing DARPA's own ARPANET in the 1990s (Roland and Shiman 2002).

Secondly, the discontinuation of 'Selective Availability' (SA), which had previously limited satellite navigation—specifically, the US Global Positioning System (GPS)—to military use (Clinton, 2000). As the original White House press release noted, 'the technology that makes this extraordinary technology [GPS] possible grows directly from our past research investments in basic physics, mathematics, and engineering supported from NSF, DARPA, NIST and other Federal agencies over a

period of decades' (Clinton, 2000, p. n.p.). Thus, after the subsequent failure of the SCI, DARPA committed to advancing GPS technologies—arguably one of many limitations in the earlier efforts to automate driving through the ALV project in the SCI. As the press release contends, switching off SA would 'have immediate implications in areas such as ... car navigation' (Clinton, 2000, p. n.p.). Here:

Previously a GPS-based car navigation could give the location of the vehicle to within a hundred meters. This was a problem, for example, in areas where multiple highways run in parallel, because the degraded signal made it difficult to determine which one the car was on. Terminating SA will eliminate such problems, leading to greater consumer confidence in the technology and higher adoption rates. It will also simplify the design of many systems (e.g. eliminate certain map matching software) thereby lowering their retail cost. (Clinton, 2000, p. n.p.)

Thus, the 'termination' of SA, as the press release puts it, not only created the conditions for a new market for GPS applications but for the start of this positional phase of automotive navigation.

The year 2009 onwards can be considered the 'executive' era, with automotive navigation being increasingly supported by an array of assistive technologies from social navigation apps such as Waze (Hind & Gekker, 2014) to advanced driver-assistance systems (ADAS) designed to offer advanced forms of vehicle control. This is arguably where the development of automotive navigation technologies and the development of autonomous vehicles begin to fuse, with the launch of Google's Self-Driving Car project (Harris, 2014a, 2014b). It is at this point—as Google's autonomous vehicle project gets off the ground—that journalist Adam Fisher speculates that the map might well end up being 'fully absorbed into the machine' (Fisher, 2013, p. n.p.).

With this background to automotive navigation in mind, this chapter proceeds to document what mapping work looks like in this new age. In particular, it considers how 'operational design domains' (ODDs)—a functional territory designed to establish the outer operational, legal, and economic limits of an autonomous vehicle—constitute a cartographic 'ground truth'. Yet ODDs are not the only way in which such mapping

work has been incorporated into the development of autonomous vehicles. In contrast, so-called mapless approaches to autonomous driving do not, necessarily, require the ongoing collection and synthesis of cartographic or geometric data pertaining to general driving environments, captured through cartographic technologies, but rather through sensing technologies. This mapless approach to autonomous driving is driven by the need for interoperability (Hind, 2023), through which geometric and semantic data must *impress* themselves on, and *express* themselves in, ML models, in the way that Matteo Pasquinelli and Vladan Joler understand machine learning as an ‘instrument of knowledge magnification’ (Pasquinelli & Joler, 2020, p. 1263). In such cases the map neither disappears nor is silently subsumed into the machine, but operates as a ghostly presence within it.

Operational Domains, Viable Terrains

For many autonomous vehicle firms, establishing an ODD is critical. ODDs constitute the outer operational limits of an autonomous vehicle. They can be considered the ‘global map’ on which the autonomous vehicle must drive. Any terrain outside of the ODD is entirely out of bounds, with the vehicle unlikely to be able to traverse an environment it lacks the basic coordinates for. In the words of Uber ATG—Uber’s now defunct autonomous vehicle division:

Before beginning any self-driving testing we establish the ODD. The ODD describes the specific conditions under which the self-driving system is intended to function, including where and when the system is designed to operate. The parameterization is not only designed to address the performance of the base vehicle platform but also system level capabilities, environmental scenarios, and appropriate self-driving system responses. (Uber ATG, 2019, p. 29)

An ODD can be considered a form of ‘ground truthing’ in itself, different from the strict machine learning definition of the term, but still critical to the development of autonomous vehicles. Without this global

map (the ground truth) the vehicle has no ultimate sense of where it is driving, even before the addition of other road users. Here, Uber ATG conceive of such work as a ‘parameterization’, that is as a process of establishing operational parameters in which their test vehicles must operate. Alongside referring to the operation of the vehicle, parameterization has a mathematical definition. In such cases, parameterization involves the process of training an ML model to perform in a particular manner, through refining and optimizing the way it interprets and acts on the data it is fed (Mackenzie, 2017). In these cases, Uber ATG might be said to have been demonstrating a kind of ‘multi-scalar’ sensibility, engaging in parameterization both at the level of the model and the operational domain. Indeed, that the practice of model parameterization—typically a process of adjusting ‘weights’ and ‘biases’—itself weighs on the parameterization of the ODD. In effect, therefore, that ODD parameterization involves the assessment of different weights attached to the input data used to determine a viable (profitable) ODD for Uber ATG.

In a second safety report released in 2020, Uber ATG outline a three-stage process involved in designing an operational domain: ‘identify, characterize, and constrain’ (Uber ATG, 2020a, p. 22).¹ At the time a division of the ride-hailing service, Uber, Uber ATG thus designed the ODD with future ride-hailing services in mind, stating:

We begin by identifying specific geographies where we would like to ultimately deploy self-driving vehicles on the Uber network by taking into consideration a number of factors, including the regulatory environment and areas where we can extend our network’s reach to better serve riders. (Uber ATG, 2020a, p. 23)

The setting of the ODD is an economic decision, driven by the demands and intentions of the parent company (Uber) its design serves. There is no pure, scientific ‘ground truth’ in which the geographical

¹ Uber ATG first released a safety report in 2018, contained within documentation released during the NTSB investigation into the fatal Uber ATG crash in Tempe, Arizona, in March 2018 (Uber ATG, 2019). The original 2018 safety report contains 70 pages, compared to a total of 89 in the updated 2020 version. A blog post announced the release of the second safety report (Uber ATG, 2020b).

extent of a specific area is mapped and determined viable from a strict safety or operational perspective. Uber ATG further make this calculation based on ‘using data from sources such as Uber’s existing lines of business [i.e. taxi rides], Uber ATG’s internal domain characterization process, and information layers of our high-definition maps’ (Uber ATG, 2020a, p. 23). In this, whilst an ODD can be considered a complete geographical area with perimeter coordinates, an ODD in Uber ATG’s case is built from route data upwards, that is that a viable ODD is based on rider data derived from Uber’s ride-hailing service. In this, a dense, entangled ride-hailing network would arguably provide a good basis for the setting of Uber ATG’s ODD. In other words, the identification of an ODD not only functions as a form of market research but more precisely as a crystallization of the *search* for viable future markets in themselves (Fig. 2.2).

The second stage involves characterizing the ODD. According to Uber ATG, this involves five separate tasks including, ‘leveraging externally-sourced data’, ‘driving the area manually’, ‘adding data tags’, ‘synthesizing the tagged data’, and ‘creating representative simulation and track tests’ (Uber ATG, 2020a, p. 24). In short, once an ODD has been identified, Uber ATG need to actually map the area in question, adding a greater level of cartographic depth and understanding, through repeated manual drives, and through specifying ‘attributes of road design’ such as ‘road

The Brain of an Uber Self-Driving Vehicle

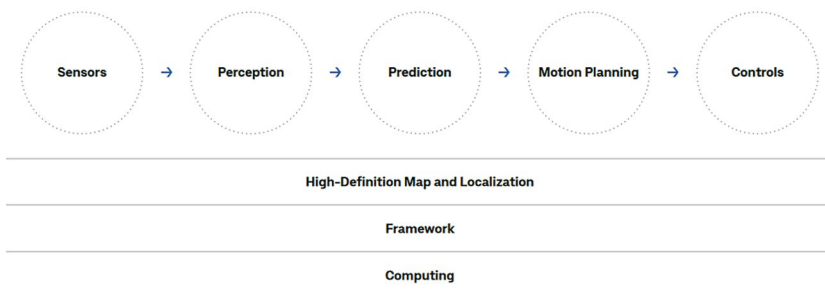


Fig. 2.2 The ‘brain’ of an Uber self-driving vehicle. (Source: Uber ATG, 2019)

geometry or curvature’ (Uber ATG, 2020a, p. 24). Characterizing the ODD involves collating the geometric data that is needed to underpin their test operations, alongside additional necessary data annotation processes. Whilst this data annotation work is integral to the building of ML models, as will be discussed in the next chapter, it is also critical to the mapping of the ODD itself. To reiterate: ‘characterizing’ the ODD is considered here as a secondary stage in the design of the operational domain, where the mapping of *geometric* features follows the processing of *semantic* features. Here, the ‘ground truth’ is derived principally from the latter, supported and validated by the former. In other words, this constitutes the formation of a ‘map-territory’ (Hind & Gekker, 2019, p. 141) designed from the semantic (social) ground up.

In the final stage before testing can begin, the ODD must be ‘constrained’ so that the prototype vehicle is prohibited from driving within, and also beyond, specific areas or spaces of the ODD (i.e. pavements or the next city), or limited from driving at particular speeds within it. For instance, as they write, ‘our self-driving system is prevented from driving in geographical areas outside of the ODD via geofencing techniques that impose a system prohibition on driving across the geofence’ (Uber ATG, 2020a, p. 25). The same techniques can be used to impose restrictions on how the vehicle drives in particular areas, ‘at the lane level based on a set of configurable ODD elements, e.g., road speed, road type, and traffic control devices’ (Uber ATG, 2020a, p. 25).

Uber ATG also write that whilst these constrain/prohibit techniques are ordinarily enacted prior to testing, as part of establishing the ODD, they are also enacted during tests themselves by employees known as ‘Mission Specialists’.² These workers are ‘trained on governing ODD, and are prepared to take manual control of the vehicle when presented with a scenario or conditions not included in the relevant ODD’ (Uber ATG, 2020a, p. 25). In short, that these workers help to prevent the vehicle from going rogue when ‘off-map’ phenomena present themselves. In such instances, the information generated by such an event ‘then initiates a process by which a live operational constraint or crew notification may be

² In other material, Mission Specialists are more plainly referred to as ‘vehicle operators’ (VOs), see Hind (2022).

created and distributed to the [autonomous vehicle] fleet with an appropriate solution, such as precluding certain future encounters with the location' (Uber ATG, 2020a, p. 25). Here, the suggestion is that each off-map object, phenomena, or location is ideally then incorporated into the ODD or definitively excluded. In each case, the heretofore unknown feature becomes knowable, ready to be responded to in future cases, by all vehicles.

As Denis Wood has considered, there has always been a politics to mapping specific automotive environments: 'that the choice to map Tibet as Chinese reveals a political attitude is something many will readily concede, but *all* choices are political and it is no less revealing to choose to map *highways*, for this too is a value' (Wood, 2010, p. 77, authors' emphasis). In his reading of a public transportation guide in the US state of North Carolina, in which he suggests the 'reek of special assistance is like sweat' (Wood, 2010, p. 77), Wood writes that 'there is nothing of this tone on the highway map' (Wood, 2010, p. 77). The highway map does not need to be requested, 'it, after all, is ... *a natural function of the state*' (Wood, 2010, p. 77, authors' emphasis):

Everything conspires to this end of naturalizing the highway map (even the map of public transportation), of making the decision to produce such a map seem less a decision and more a gesture of instinct, of making the map's cultural, its historical, its political imperatives transparent: you see through them, and there is only the map, innocent, of nature, of the world as she really is. (Wood, 2010, p. 77)

The development of autonomous driving both builds on, and somewhat departs from, this 'instinctive' impulse. Firstly, that highways are considered a 'pure' unimpeded driving environment already cleared of the messiness of the world, vulnerable road users, and unpredictable others. As a result, the highway constitutes a quasi-laboratory space for the testing of autonomous vehicles, with few(er) obstacles to classify and negotiate, and long stretches without the need, essentially, to make the kinds of decisions autonomous vehicles struggle with (Hind, 2019).

Yet, secondly, as a result of this, the highway also represents only a small part of the driving experience for many drivers who, on most days,

do not have to negotiate a four-lane highway or drive from city to city. Instead, the automobile is used to drive from the suburbs into a small town, to an office car park, or an out-of-town supermarket. For many (but not all) people, this ordinarily does not involve driving on a highway but on other types of roads where there is a greater mix of road users, a greater number of possible obstacles to negotiate, and therefore a far more complex array of risks any autonomous vehicle must in some sense calculate.

Thus, whilst mapping highways might indeed be a ‘natural function of the state’, as Wood (2010, p. 77) writes, it isn’t necessarily the natural function of the autonomous vehicle company. Instead, determining ODDs becomes the natural function of the autonomous vehicle firm, making economic calculations concerning questions of financial ‘viability’ and market ‘sustainability’. Highways might well work their way into these calculations, but, ultimately, it is a question of whether the inclusion of highways into the ODD of an autonomous vehicle firm succeeds in passing a viability test first.

For Tesla, manufacturer of a problematic ‘autopilot’ feature within their vehicles, highways have constituted a principle viable typology (Hind, 2019). For Waymo, the more intricate lattice of (sub)urban roads have constituted their ‘primary test environment’, despite (or perhaps because of) exhibiting ‘a rather more unruly set of social phenomena’ (Hind, 2019, pp. 411–412). For Uber ATG, in the business of developing an autonomous ride-hailing service, the case is built on the basis of Uber’s existing operational markets. Some places are more viable, and more valuable, than others. Mapping these ‘viable’ and valuable spaces, therefore, is a critical part of the autonomous vehicle process.

Semantic Semantics

As Alex Gekker and I have considered, other autonomous vehicle players also recognize the critical importance of establishing ‘detailed semantic understanding of how people *usually* drive’ (Hind & Gekker, 2024, p. 3715, emphasis added), in addition to collecting details on the external laws that permit them to drive. Yet, collecting this semantic-level data is

somewhat novel from a cartographic perspective, new also from an automotive perspective, although similar in type to behavioural data broadly collected through digital devices and platforms. In other words, the desire to collect ‘detailed semantic understanding of how people usually drive’ is connected to the appetite for the collection of semantic-level data on all kinds of everyday human behaviour, from exercising to shopping. What is different here, I argue, is that this data has *immediate operational valuable* rather than indirect or latent circulatory value, that is that it is being collected and used by autonomous vehicle firms precisely because it offers an understanding of ordinary driver behaviour that can be emulated in a resultant autonomous vehicle. That is, rather than being used for commercial data-based advertising opportunities or user surveillance (Gekker & Hind, 2020).

For firms such as Mobileye, this involves developing something called ‘Road Experience Management’ (REM), a process through which the firm extracts semantic data from real drivers. As an established ADAS manufacturer—devices that attach to the dashboard of a vehicle, acting as a warning system for drivers—it has been able to amass huge volumes of data from users to build their autonomous vehicle platform. In the words of their CEO, Amnon Shashua, these largely unsuspecting ADAS users comprise ‘millions of harvesting agents’ (Mobileye, 2021, p. n.p.) around the world, providing Mobileye with a unique position in the industry of being able to capture data from drivers irrespective of car they drive, as Hind and Gekker (2024) contend.

Yet, what makes the autonomous vehicle start-up Waabi different from Mobileye is that it can’t rely on millions of pre-existing users of in-house devices, already in the cars of drivers the world over. Neither can it rely on a fleet of vehicles to generate such data—Waabi doesn’t possess the kind of capital available to big tech competitors like Waymo and (previously) Uber ATG, that makes running a fleet of test vehicles feasible. As a result, it has had to think creatively about how it resolves certain questions about the acquisition of both geometric and semantic data. Their solution? To extract geometric and semantic data directly from lidar itself (Casas et al., 2021), skipping both an ODD phase where geometric data is collected independently from semantic data and any pre-processing of

lidar data in advance of extracting semantic-level insights. It is this I turn to next.

(The Promise of) Mapless-ness

In the case of Waabi, engineers have developed an alternative to ‘high-definition’ (HD) maps called MP3: an ‘end-to-end approach to mapless driving’ (Casas et al., 2021, p. 1), where MP3 stands for ‘map, perceive, predict and plan’ (Casas et al., 2021, p. 1).³ This approach differs from HD map techniques containing ‘rich semantic information necessary for driving’ (Casas et al., 2021, p. 1), such as those by Mobileye, Uber ATG, or mapping firm HERE. Following the MP3 approach, ‘scene representations’ that feed into motion planning are driven by the live extraction of ‘meaningful geometric and semantic features from ... raw sensor data’ (Casas et al., 2021, p. 1). In short, that Waabi envisage both mapping and sensing as taking place in tandem and in real-time; rather than mapping taking place prior to any live sensing, as is common in other HD map approaches. In this sense, the sequential ordering of mapping and sensing operations—reflected in the ordering of the chapters in this book—is only correct for some operators in the domain. For Waabi, there is no sequential ordering, only parallel ingestion (Fig. 2.3).




Fig. 2.3 ‘Map, perceive, predict and plan’: The architecture of Waabi’s MP3 model. (Source: Casas et al., 2021)

³The original technical explication of MP3 by Sergio Casas and colleagues (Casas et al., 2021) was written whilst all three authors worked for Uber ATG. All three authors now work for Waabi whilst two hold academic positions at the University of Toronto (Urtasun and Casas).

This is why Waabi considers such an approach to be essentially ‘mapless’: no HD, granular maps of pre-existing road infrastructure, or semantic levels are produced in advance of a vehicle being autonomously driven. As Casas et al. (2021, p. 1) contend, ‘heavy reliance on HD maps introduces very demanding requirements for the localization system, which needs to work at all times with centimeter-level accuracy’ without fear of an accident. In such cases, the vehicle in question senses the environment and calculates possible path trajectories in relation to already mapped phenomena. Autonomous vehicle trials, the likes of which have been seen on the streets of Pittsburgh and Tempe, Arizona, are typically as interested in establishing ODDs as they are in testing, improving, sensing, and decision-making itself.

Why is Waabi’s foregrounding of a ‘mapless’ approach to autonomous driving important to understand? Firstly, that despite being explicit about such, they aren’t necessarily the only operator to offer such an approach. As I will examine in Chap. 7, open-source manufacturer Comma utilizes an ‘end-to-end approach’ (E2E), meaning that their ML models are driven by vehicle data they ingest (from devices purchased by users) and nothing else. Whilst a mapless approach has arguably been less common in the industry, it avoids the endless rule-setting and codification that has plagued other autonomous vehicle set-ups.

What the mapless approach offers is the possibility to streamline operations. The chapters in this book, loosely, reflect the modular operations that are executed when autonomous vehicles make decisions. What Waabi want to effectively do is to cut out one such module: the part where they must exhaustively map an ODD. Not only does this process naturally take time, labour, and considerable resources, but it also must be repeated iteratively to ensure the ODD matches external reality. Any slippage between these two—a road made one-way only, a cycle lane installed—and the autonomous vehicle relying on an outdated ODD obviously becomes a huge danger to other road users. Waabi’s mapless approach, in theory, skips this step and avoids having to deploy vehicles to establish the geometrics of an area prior to operation. Instead, as articulated before, it folds this work into subsequent phases of operation, namely the ML model training phase.

Urtasun justifies this approach—which is not without considerable risk—by arguing that one of the main hurdles in the development of autonomous vehicles is a modular approach to system design favoured in the tech industry. Whilst I will return to this in Chap. 4 in relation to sensing and sensor work, Urtasun argues that most technical solutions to extant developmental problems ordinarily involve the formation of additional groups of engineers, and the construction of additional modules, to fix the problem (Urtasun, 2021). As she explains:

If you think about the traditional software stack, there are a couple of advantages. But also, there are some issues. The system is easily interpretable, and it's very easy to incorporate prior knowledge. Engineers can, for example, tune the cost of the motion planner in order to reflect their intuitions about the task of driving. However, if you look at how this is typically deployed or worked on in industry, you're going to have teams of hundreds of people – in some of the companies more than a thousand – working on this piece of software. So, you end up with teams of work in silos, in small pieces, and there is no really holistic view into how to solve these tasks as a single system, which is extremely important to do, for such a complicated task as self-driving. (Urtasun, 2021, p. n.p.)

Thus, the unwieldy, modular, work involved in developing autonomous vehicles comes to weigh heavily on the speed, and productivity, of the outputs being generated:

Furthermore, typically the software stacks don't look as pretty as a few lambda functions ... instead it's an extremely complicated system, where every time there is an issue on the road basically you tackle one more model and one specific thing, and it looks more like an if-then statement than anything else. Also, the fact that there is no ability to train this end to end, you end up basically with many little interfaces between the different models, and as a consequence if you have a mistake, it's very hard to actually ... correct this mistake. Instead, you have this cascading issue as you go through the software stack. Engineers typically are pretty unhappy because developer productivity is very low, because there is no automation, basically, you have to tune one thing at a time, and every time that you make a

change in one of the models, you will have to change everything else. (Urtasun, 2021, p. n.p.)⁴

The alternative to this, as Urtasun contends, is an E2E or ‘mapless’ approach: ‘sensors in, driving commands out’ as Urtasun (2021) puts it:

The advantage of this is super simple. A few lines of code will do it ... the other advantage is that you’re going to train the system for the end task of driving. Now, the difficulty of this problem is that really there is no interpretability. In the system if there is something wrong, you’re going to have an issue explaining why that is the case. Also, it’s very difficult to incorporate prior knowledge, and the assumption is that you’re going to observe phenomena, and you’re going to have data for all the things that might happen, and then basically, you’ll learn from it. But as we know, it’s potentially unethical to collect data where you have near misses or accidents. So, this is not a solution either. (Urtasun, 2021, p. n.p.)

Yet, Urtasun isn’t especially happy with one key aspect of this strictly E2E approach: that isolating and understanding any inherent issues are made harder by developing a single ‘sensors in, driving commands out’ model. Thus, she offers a third solution, a kind of qualified E2E—still mapless—approach, that attempts to resolve some of these interpretability issues encountered with a single ML model:

So, the type of technology we’ve been developing ... is something that incorporates the advantages of these two approaches, without incorporating the disadvantages. In particular, it’s going to be one model, and you’re going to be able to incorporate prior knowledge, and it’s going to be interpretable, and you’re going to be able to explain why the system decides to do a particular manoeuvre. However, it’s going to be an end-to-end trainer, so you have all these advantages of being able to have more complex functions, even for example, when you’re doing the motion planner model, you can actually get access to the raw data, which is potentially very, very, important – so you don’t have these cascading mistakes. (Urtasun, 2021, p. n.p.)

⁴Lambda or ‘anonymous functions’ are algorithmic processes used for simple expressions and as such can be used variously in computer programming.

Urtasun's articulation of Waabi's approach to autonomous driving can be understood as a drive for interoperability (Hind, 2023), in which raw cartographic data in both geometric and semantic form are not only vital for the technical work of model building but also for the *interpretability* of such model building. Whereas a modular approach leads to infinite 'little interfaces' and 'cascading errors', as Urtasun explains, a single model built direct from mapping and sensing data to driving commands only serves to obfuscate, and complicate, proceedings. If things go wrong, where can one find the error? An E2E trainer—ideally, if not actually—that is both (a) holistic and (b) explainable is, in Waabi's world, the best of both these worlds. It is thus a promise: of mapless mapping where mapping and sensing data, geometric and semantic data, must impress themselves on, and express themselves in, ML models. Getting the model to fully express those impressions to engineers when things go wrong, is precisely the challenge.

'As an instrument of knowledge', Pasquinelli and Joler (2020, p. 1265, authors' emphasis) write, 'machine learning is composed of an object to be observed (*training dataset*), an instrument of observation (*learning algorithm*) and a final representation (*statistical model*)'. Urtasun's efforts are an attempt to fine-tune the relationship between these three components: the training dataset, the learning algorithm(s), and the statistical model(s). Indeed, where things are even trickier for Urtasun and her team, is that there are necessarily multiple 'final representations', each—once connected together, fused and interoperable—liable to generate 'cascading errors'. Yet, Urtasun's holistic quest, embodied in the architecture of Waabi's MP3 model, still runs the risk of delimiting a necessary kind of explainability for which the 'many little interfaces' are surprisingly good at surfacing. Here, the streamlining of operations always runs in tension with interpretability and accessibility of the model. To run with the pipeline metaphor: it's like a sewage system without an inspection hole.

Following Pasquinelli and Joler's analogy, 'the information flow of machine learning is like a light beam that is projected by the training data, compressed by the algorithm and diffracted towards the world by the lens of the statistical model' (Pasquinelli & Joler, 2020, p. 1265). This is why both Urtasun (2021) and Casas et al. (2021) emphasize the

importance of ‘injecting prior knowledge’ into the models being built: the original projections must shine brightly and be directed accurately even before they are compressed and diffracted. Yet ‘sensors in, driving commands out’ isn’t quite as straightforward as would be hoped, and yet as Urtasun (2021) suggests, ‘it’s very difficult to incorporate prior knowledge’ if there are all sorts of necessary—practical, ethical—qualifications to incorporating such.

It is worth saying a little more on Urtasun (2021) and Casas et al.’s (Casas et al., 2021) proposed mapless solution to the modularity of model development on one side and a lack of interpretability on the other side. As Urtasun’s (Urtasun, 2021) ‘sensors in, driving commands out’ quip suggests, the only data that enters the model is raw sensor data. It is from this lidar data that both geometric *and* semantic features are extracted. This differs from a map-based approach because geometric data is ordinarily collected through the establishment of an ODD, not extracted from lidar data itself. From here, ‘voxelized’ lidar data forms the input for a backbone network. Next, aforementioned ‘scene representations’ are generated through two processes: the generation of (a) an ‘online map’ meant to encode the fabled ‘prior knowledge’ of human drivers and (b) a ‘dynamic occupancy field’, recording the position and velocity of other ‘dynamic objects’ on the road, that is other road users. From here the online map and dynamic occupancy field—both only generated from lidar data—are fed into a ‘motion planner’ used to set ‘kinematically-feasible trajectories’ (Casas et al., 2021, p. 5) for the autonomous vehicle itself.

Yet, even a mapless vision contains maps. Firstly, the Waabi team explicitly acknowledge the role of off-the-shelf ‘course road network’ maps (Casas et al., 2021, p. 1) for final vehicle routing purposes, for which they don’t—sensibly—propose a mapless alternative. Yet, as seen from the above description of the MP3 model, the engineers also talk of an ‘online map’ integral to the representation of the ‘swarming social reality’ (Hind, 2019, p. 412) within which the autonomous vehicle must inhabit. Here, whilst the engineers don’t qualify the ‘mapiness’ of their mapless approach for a second time, it becomes obvious that—somewhere down the line—maps find a way of inserting themselves back into the equation, best equipped (sometimes, if not all the time) to render and

represent phenomena in what the technical paper refers to as ‘BEV’: engineering shorthand for ‘bird’s eye view’.

To return to the original critique of the ODD-reliant, modular, ‘traditional software stack’ approach Urtasun (2021) makes, the main problem is compartmentalization: of the work being performed, of the errors surfaced, of the solutions found, of the overall decision-making system designed, and ultimately, of the exasperation felt. HD maps are the embodiment of this compartmentalization: labour-intensive to create, and unwieldy to manage and update. Yet as I’ve unpacked in detail before (Hind, 2023), developing technical solutions to extant problems in autonomous driving always consists of computational trade-offs, often between speed and accuracy. In the case of Waabi’s MP3 model, they must live with a lack of so-called ‘HD maps’ from start to finish, constantly finding alternatives and solutions to that which they provide. This forms a kind of interoperational ‘drag’ on proceedings all the way through, from being unable to exploit lane geometry to lane sequency. This drag ultimately constitutes an ongoing cost or tariff on the development process, as well as a looming reality not yet touched upon: that an ODD-dependent autonomous vehicle system is seen by regulators as being the safest approach to autonomous driving. With ‘mapless-ness’ either comes fearlessness or recklessness, another tension at the heart of the pursuit to automate driving.

Conclusion

In this chapter I have chronicled the efforts of two firms involved in mapping the spaces of autonomous driving. In the first instance, I have discussed the work of Uber ATG—previously Uber’s autonomous vehicle division—to develop so-called operational design domains or ODDs. Considered a kind of ‘ground truth’, ODDs are meant to establish the outer operational limits of autonomous vehicles undergoing testing. As I have argued, however, Uber ATG’s development of ODDs is not a ‘pure’ scientific pursuit, but one guided principally by the search for viable future markets. Here, the objective of developing autonomous ride-hailing services doesn’t just seep into the technical work of determining

the ODD but is the driving force behind it. This search for viable terrains, thus, consequently bends all other features towards this goal—from the leveraging of external data to the precise mapping of roadside curbs. All this mapping work—from data annotation to data synthesis—crystallizes the search for viable markets.

In the second instance, I have discussed the work of a different company altogether—an autonomous vehicle start-up called Waabi. Founded by Uber ATG's former Chief Scientist, Raquel Urtasun, Waabi doesn't have what others like Uber ATG or Mobileye possess: access to huge volumes of pre-existing data. In order to deal with these and other deficiencies, Waabi offers a vision of a 'mapless' world, where autonomous vehicles ingest, and operationalize, sensor data on the fly. Designed to side-step the issue of a lack of pre-existing data, Waabi's approach is also designed to tackle another problem: that building an autonomous vehicle system is a 'modular' affair, destined to generate cascading errors. Removing a mapping phase entirely from this process is Urtasun's solution bundling mapping and sensing phases together as one. The result is an innovation: a kind of 'mapless mapping' generated through cartographic *impressions* destined for their *expression* in resultant ML models.

References

- Akerman, J. (2006). Twentieth-century American road maps and the making of a national motorized space. In J. Akerman (Ed.), *Cartographies of travel and navigation* (pp. 151–206). The University of Chicago Press. <https://doi.org/10.7208/chicago/9780226010786.003.0005>
- Casas, S., Sadat, A., & Urtasun, R. (2021). MP3: A unified model to map, perceive, predict and plan. Preprint. <https://doi.org/10.48550/arXiv.2101.06806>
- Clinton, B. (2000, May 1). *President Clinton: Improving the civilian Global Positioning System (GPS)*. The White House: Office of the Press Secretary. https://clintonwhitehouse4.archives.gov/WH/New/html/20000501_2.html
- Dant, T. (2004). The driver-car. *Theory, Culture and Society*, 21(4–5), 61–79. <https://doi.org/10.1177/0263276404046061>
- Fisher, A. (2013, December 11). Google's road map to global domination. *The New York Times Magazine*. <https://www.nytimes.com/2013/12/15/magazine/googles-plan-for-global-domination-dont-ask-why-ask-where.html>

- French, R. L. (2006). Maps on wheels: The evolution of automobile navigation. In J. Akerman (Ed.), *Cartographies of travel and navigation* (pp. 260–290). The University of Chicago Press. <https://doi.org/10.7208/chicago/9780226010786.003.0007>
- Gekker, A., & Hind, S. (2020). Infrastructural surveillance. *New Media & Society*, 22(8), 1414–1436. <https://doi.org/10.1177/1461444819879426>
- Harley, J. B. (1989). Deconstructing the map. *Cartographica*, 26(2), 1–20. <http://hdl.handle.net/2027/spo.4761530.0003.008>
- Harris, M. (2014a, September 10). How Google's autonomous car passed the first US state self-driving test. *IEEE Spectrum*. <https://spectrum.ieee.org/how-googles-autonomous-car-passed-the-first-us-state-selfdriving-test>
- Harris, M. (2014b, November 19). The unknown start-up that built Google's first self-driving car. *IEEE Spectrum*. <https://spectrum.ieee.org/the-unknown-startup-that-built-googles-first-selfdriving-car>
- Hind, S. (2019). Digital navigation and the driving-machine: Supervision, calculation, optimization, and recognition. *Mobilities*, 14(4), 401–417. <https://doi.org/10.1080/17450101.2019.1569581>
- Hind, S. (2020). Between capture and addition: The ontogenesis of cartographic calculation. *Political Geography*, 78, 1–9. <https://doi.org/10.1016/j.polgeo.2020.102147>
- Hind, S. (2022). Machinic sensemaking in the streets: More-than-lidar in autonomous vehicles. In G. Rose (Ed.) *Seeing the city digitally: Processing urban space and time* (pp. 57–80). Amsterdam: Amsterdam University Press. <https://doi.org/10.2307/j.ctv2j6xrs3.6>
- Hind, S. (2023). Sensor work: Enabling the inter-operation of autonomous vehicles. *Journal of Computer-Mediated Communication*, 28(5), 1–10. <https://doi.org/10.1093/jcmc/zmad014>
- Hind, S., & Gekker, A. (2014). 'Outsmarting traffic, together': Driving as social navigation. *Exchanges: The Warwick Research Journal*, 1(2), 1–17. <https://doi.org/10.31273/eirj.v1i2.84>
- Hind, S., & Gekker, A. (2019). On autopilot: Towards a flat ontology of vehicular navigation. In C. Lukinbeal, L. Sharp, E. Sommerlad, & A. Escher (Eds.), *Media's mapping impulse* (pp. 141–160). Franz Steiner Verlag. <https://doi.org/10.25162/9783515124256>
- Hind, S., & Gekker, A. (2024). Automotive parasitism: Examining Mobileye's 'car-agnostic' platformisation. *New Media & Society*, 26(7), 3707–3727. <https://doi.org/10.1177/14614448221104209>

- Kitchin, R., & Dodge, M. (2007). Rethinking maps. *Progress in Human Geography*, 31(3), 331–344. <https://doi.org/10.1177/0309132507077082>
- Mackenzie, A. (2017). *Machine learners: Archaeology of a data practice*. MIT Press.
- Mobileye. (2021). *Road experience management*. Mobileye. Retrieved October 1, 2023, from <https://www.mobileye.com/technology/rem/>
- Pasquinelli, M., & Joler, V. (2020). The Nooscope manifested: AI as instrument of knowledge extractivism. *AI & Society*, 36, 1263–1280. <https://doi.org/10.1007/s00146-020-01097-6>
- Roland, A., & Shiman, P. (2002). *Strategic computing: DARPA and the quest for machine intelligence*. MIT Press.
- Schulten, S. (2000). The limits of possibility: Rand McNally in American culture, 1898–1929. *Cartographic Perspectives*, 35, 7–26. <https://doi.org/10.14714/CP35.834>
- Sheller, M., & Urry, J. (2000). The city and the car. *International Journal of Urban and Regional Research*, 24(4), 737–742. <https://doi.org/10.1111/1468-2427.00276>
- Simondon, G. (2017). *On the mode of existence of technical objects*. Univocal Publishing.
- Singh, R., Hesselbein, C., Price, J., & Lynch, M. (2019). Getting ‘there’ from the every-changing ‘here’: Following digital directions. In J. Vertesi & D. Ribes (Eds.), *digitalSTS: A field guide for science & technology studies* (pp. 280–299). Princeton University Press.
- Thielmann, T. (2016). Linked photography: A praxeological analysis of augmented reality navigation in the early 20th century. In E. Gómez Cruz & A. Lehmuskallio (Eds.), *Digital photography and everyday life: Empirical studies on material visual practices* (pp. 160–185). Routledge.
- Uber ATG. (2019, November 5). *Safety report: A principled approach to safety*. NTSB. Retrieved October 1, 2023, from <https://data.nts.gov/Docket/Document/docBLOB?ID=40477737&FileExtension=.PDF&FileName=Operations%20Attachment%20-%20Uber%20Advanced%20Technologies%20Group%20A%20Principled%20Approach%20to%20Safety%202018-Master.PDF>
- Uber ATG. (2020a, August 28). *Safety Report 2.0: A principled approach to safety*. Uber ATG. Retrieved October 1, 2023, from <https://uber.app.box.com/v/UberATGSafetyReport>
- Uber ATG. (2020b, August 28). *Uber ATG releases 2020 safety report*. Uber ATG. <https://medium.com/@UberATG/uber-atg-releases-2020-safety-report-575db33f2bd7>

- Urtasun, R. (2021, June 20). CVPR'21 WAD keynote—Raquel Urtasun, Waabi/University of Toronto. *YouTube*. Retrieved October 1, 2023, from <https://www.youtube.com/watch?v=PSZ2Px9PrHg>
- Wilken, R., & Thomas, J. (2019). Maps and the autonomous vehicle as communication platform. *International Journal of Communication*, 13, 2703–2727. <https://ijoc.org/index.php/ijoc/article/view/8450>
- Wood, D. (2010). *Rethinking the power of maps*. The Guilford Press.



3

Training Decisions: Ground-Truthing the Interesting

This chapter is concerned with how training datasets are composed and ‘put to work’. In this case, training datasets do not exist on their own as discrete entities. Instead, training datasets—as their names suggest—are necessary components for training machine learning (ML) models. Training datasets, thus, are understood within the ML community as a so-called ground truth, representing the known, and foundational, statistical terrain upon which ML algorithms learn to operate. Once such a model has been trained, using a particular training dataset, it is ostensibly ready to be applied on new data that may or may not share a resemblance with the training dataset it has been trained on. As a result, practitioners training ML models look to utilize a training dataset that can be said to be both *large enough* to incorporate enough, repeat variations of situations relevant to the application context, and *diverse enough* to incorporate a variety of situations relevant to the same context.

In the case of autonomous driving, this ordinarily translates into (a) a volume of familiar road ‘scenes’ and environments, from highways to urban junctions and (b) a variety of possible situations a vehicle is likely to encounter at any given time. In summary, that training datasets in autonomous driving are valued for their size and diversity. Whilst this is something it shares with other domains in which ML is being applied,

these situations are naturally unique to a world of driving, as opposed to those might occur at an international border or in the supermarket. It is this interleaving of the general and the specific that this chapter will consider.

In this, questions of ‘overfitting’ (where an ML model performs well on a training dataset but poorly on new data) and ‘underfitting’ (where an ML model performs poorly on both training data and new data) underpin an understanding of how an ML model performs in relation to training data. Both situations are undesirable and can be resolved in numerous ways, in order to achieve a desired ‘fit’.

In addition, a suite of technical considerations such as sensor type (lidar, camera), sensor combination (lidar and camera, lidar and radar), and computational hardware, as well as some understanding of how the training dataset has been aggregated, combine to give value to the training dataset itself. However, as this chapter will consider, training datasets are ordinarily considered as ‘given’ inputs when ML work is being performed, only ever the ground (truth) on which the models are produced. In this, not only may they be rarely questioned, but as particular datasets are used, and re-used, their validity or their ground-truthiness hardens. This chapter, therefore, will consider two particular training datasets within the autonomous driving world that have ‘hardened’ to different degrees: the KITTI Vision Benchmark Suite, first released in 2012, and Waymo Open Dataset, publicly released in 2019.

In recent years, two sets of authors have written similar statements regarding AI and the use of training datasets. Firstly, Nick Srnicek has argued that thanks to the spread of the ‘data-centric business model’ (Srnicek, 2022, p. 248) led by platform firms such as Google/Alphabet, ‘the problem of how to bootstrap [AI systems] from no data is, simply put, increasingly less of a problem’ (Srnicek, 2022, p. 248). Following this argument, collecting data for the purposes of training ML models in a range of cases is now considered less of a burden than before. Indeed, that with KITTI, Waymo Open Dataset (mentioned by Srnicek), and many others such as Motional’s NuScenes (fully released in 2019) and Argo’s Argoverse (2019), the autonomous driving domain is now replete with training datasets.

Secondly, Kate Crawford and Trevor Paglen have suggested that thanks to decades of machine vision research, ‘challenges such as object detection and facial recognition have been largely solved’ (Crawford & Paglen, 2019, p. n.p.). At least, that is, according to the ‘art of inevitability [that] recurs in many AI narratives, where it is assumed that ongoing technical improvements will resolve all problems and limitations’ (Crawford & Paglen, 2019, p. n.p.). Following this argument, AI practitioners believe that the underlying, core aspects of machine vision tasks like object detection have largely been figured out, and that with each successive year, technical progress—largely incremental in form—will continue to be made.

Srnicek’s restating of this typical claim shouldn’t, however, be read as an uncritical adoption of one such AI narrative (‘we have all the data we need for machine vision’) but an identification of the stated shift in interest, emphasis, and focus, by AI firms towards computation and labour. Here, with data collection and aggregation ‘resolved’ for all intents and purposes, attention has moved on to how to use it. Crawford and Paglen’s (Crawford & Paglen, 2019, p. n.p.) question of whether ‘the challenge of getting computers to “describe what they see” will always be a problem?’ is rarely considered in a domain like autonomous driving, singularly committed to delivering incremental performance gains in object-recognition processes.

The chapter proceeds by first establishing the status of training datasets as objects of critical inquiry, before considering the role of classification and ground-truthing in their use in training ML models. The chapter then moves on to compare the KITTI Vision Benchmark Suite and Waymo Open Dataset, two training datasets each representative of a different era—benchmark, incremental—of autonomous driving. In this, the chapter examines the significance of ‘interesting-ness’ as a sought-after quality by machine vision researchers, a quest as seemingly endless as it is challenging.

Training Datasets

As Nanna Bonde Thylstrup writes, up until recently, datasets—training or otherwise—‘have been allowed to lead a quiet existence as unassuming data feeding into fantastic new artificial intelligence (AI) inventions’ (Thylstrup, 2022, p. 656). In particular, Thylstrup suggests that datasets used for AI purposes have largely, and only, been considered as ‘operational instruments’ (Thylstrup, 2022, p. 656) rather than as critical, epistemological spaces. In this section then, I attempt to balance an interest in training datasets as operational instruments and the work performed on, and with them, with offering an account of training datasets used within the domain of autonomous driving that understands them as partial, constructed, synthetic objects.

To surface the training dataset as a contestable space, Thylstrup makes a connection to critical archival studies, which has sought to challenge ‘traditional perceptions of archives as holders of truth’ (Thylstrup, 2022, p. 657), drawing on how archives can be understood as ‘iterative spaces in which we repeat, rehearse, re-encounter and re-member the past as present’ (Thylstrup, 2022, p. 658). Whilst this is offered as a critical, and therefore necessarily counter-, understanding of how archives operate, at the stage of their composition, training datasets are often encountered in this manner also. In the labelling of training data for autonomous driving, as I will discuss later, annotation work is itself defined by a certain ‘recursivity’ in which annotators return to their annotation work in different ways, for instance, to evaluate the quality or accuracy of their work, or to re-establish, or to redefine, annotation protocols. As Thylstrup later acknowledges, research disciplines are just as engaged in ‘calls for data set context’ (Thylstrup, 2022, p. 661) as those involved in activist work, and these ‘organizational and conceptual reconfigurations [that] have transformed the value and mobility of data sets’ (Thylstrup, 2022, p. 661).

As Thylstrup thus defines it, a critical reading of training datasets in AI work:

Collectively points to data sets as a central and relational object of concern in machine learning, offering close readings of how data sets are collected,

organised, distributed and deployed, as well as explorations of their aesthetics and affective compositions. (Thylstrup, 2022, p. 659)

What is important in this case, then, is how training datasets are put to work along the autonomous vehicle pipeline, settled between previous stages of cartographic, semantic, and sensor data collection and the development of ML models themselves. Here, training datasets are the lynchpin or the interface between data collection, on one side, and AI deployment on the other. Without training data being compiled and annotated, machine learning cannot proceed. As such, and as I explore throughout the book, I understand the work performed on such data, and the work to certify it as training data, as necessary to securing the interoperability of autonomous vehicles.

Much of the work undertaken in this vein, therefore, is principally focused on smoothing, and formatting data, so that their provenance can be assured but ultimately forgotten about or backgrounded. Interoperability necessarily depends on this stated givenness, but that givenness requires knowledge of their construction to be erased. Training datasets, ultimately, are meant to become given, and this chapter will consider how this givenness is secured in the context of certain foundational datasets such as KITTI, and being sought by new ones, such as Waymo Open Datasets.

This specific sense of interoperability I refer to here, of the interoperability required of data as it flows through an ML pipeline, is not, necessarily, the same definition as used by others. As Thylstrup also writes, ‘as data sets become more interoperable and more easily shared, partitioned and modified, they also increasingly challenge complete removal’ (Thylstrup, 2022, p. 662). In this, Thylstrup understands interoperability as the interchangeability and ‘swap-ability’ of datasets, able to be plugged into a different ML process and context without hassle or issue. This is what one might call a *horizontal* interoperability in which the same training data can, hypothetically at least, be used in different AI pipelines or, perhaps, even whole domains altogether. In the case of the KITTI training dataset to be discussed, this horizontal interoperability has been designed-in from the beginning to the extent that the training data is designed to be used in any ML workflow, standardized format, and

industry-standard software. This is a form of interoperability that differs from kind I work with here, of *vertical* or linear interoperability in which training datasets, as well as all other modular operations, connect to other modules concerning mapping, sensing, motion planning, and the like.

Classifying, Annotating

Training datasets are important for machine learning because they establish the ‘ground’ on which it operates, whilst enabling a form of interoperability between dataset and pipeline. But the training dataset attains its own power through its ability to classify and categorize. ‘If we understand machine learning as a data practice’, as Adrian Mackenzie (2017, p. 9) implores us to, ‘then differences associated with machine learners in the production of knowledge should be a focus of attention’ (Mackenzie, 2017, p. 9). These differences ‘are an operative concern’, then, ‘because many machine learners classify things’ (Mackenzie, 2017, p. 9), such that ML processes can be defined as ‘classifiers’ in themselves for the work done to sort and differentiate.

Bowker and Leigh-Star previously suggested that ‘few have looked at the creation and maintenance of complex classifications as a kind of work practice’ (Bowker & Leigh Star, 1999, p. 5), despite the moral and cultural forces that not only power them but bake them into ‘the modern information technology world’ (Bowker & Leigh Star, 1999, p. 5). Thus, as Bechmann and Bowker (2019) have considered, such classification work is integral to the seemingly autonomous characteristics of contemporary machine learners. Here, data collection, data cleaning, and model training are all considered classificatory steps in the ML process, where machine learners might simply ‘assume’ or ‘derive’ *a priori* categories from ‘institutionalized or accepted knowledges’ (Mackenzie, 2017, p. 10). More than this, and perhaps much closer to a specific understanding of how machine learning works probabilistically and iteratively, that machine learners ‘invent or find new sets of categories for ... particular purpose[s]’ (Mackenzie, 2017, p. 10). These categories, with specific applications in mind, are always therefore domain dependent.

Yet, as Crawford and Paglen contend, training datasets ‘share some common properties’ (Crawford & Paglen, 2019), regardless of context or domain. The key property they all share is that the data within them is labelled and categorized, often referred to in technical literature as being ‘annotated’ and sorted into ‘classes’. For machine vision purposes, training data annotation is a process through which the spatial extent of any given feature or object within the data itself is recorded. For 2D camera data in the context of autonomous driving, this would ordinarily consist of a set of datapoints that constitute the outline of a particular object in the immediate environment, depending on the precision of the annotation work, and the specified granularity of the annotations. By way of example, the Cityscapes Dataset, compiled by researchers at TU Darmstadt, Daimler AG, Max Planck Institute for Informatics, and TU Dresden, consists of 5000 images with ‘fine’ annotations and 20,000 images with ‘course’ annotations (Fig. 3.1) (Cordts et al., 2016).

In addition, on a level up from annotation, or a step after labelling, is the sorting of annotated objects into categories or classes. In the case of the Cityscapes Dataset, there are 30 such classes, ranging from road, sidewalk, parking, and rail track (grouped under ‘flat’) to car, truck, bus, on rails, motorcycle, bicycle, caravan, and trailer (grouped under ‘vehicle’) (Cityscapes, 2022). Each class also contains a definition delimiting what is included within each class, as opposed to another. The road class, for instance, is defined as:

Part of ground on which cars usually drive, i.e. all lanes, all directions, all streets. Including the markings on the road. Areas only delimited by markings from the main road (no texture change) are also road, e.g. bicycle lanes, roundabout lanes, or parking spaces. This label does not include curbs. (Cityscapes, 2022, p. n.p.)

Whilst the traffic sign class (grouped under ‘object’) is defined as:

Sign installed from the state/city authority, usually for information of the driver/cyclist/pedestrian in an everyday traffic scene, e.g. traffic signs, parking signs, direction signs – without their poles. No ads/commercial signs. Only the front side of a sign containing the information. The back side is



Fig. 3.1 Difference between coarse (top) and fine (bottom) annotations in the Cityscapes Datasets. (Source: Cityscapes, 2022)

static. Note that commercial signs attached to buildings become *building*, attached to poles or standing on their own become *static*. (Cityscapes, 2022, p. n.p., authors' emphasis)

What is notable in both is that the respective definitions attempt to provide an exhaustive list of all possible types of instances within each class (all lanes, all directions, traffic signs, parking signs, etc.), whilst also mentioning objects or instances that might, without a definition, have been included within each class, for example curbs or commercial signs.

In other definitions, such as the one for car, it is specified that if any one particular vehicle cannot be distinguished from another ('if the boundary between such instances cannot be clearly seen'), then 'the whole crowd/group is labeled together and annotated as group, e.g. car group' (Cityscapes, 2022, p. n.p.). Thus, whilst specific vehicles may indeed be independent objects, they might be treated as a group of objects if the scene captured cannot identify the spatial extent of each particular object.

As Jatón's (Jatón, 2021) interlocutors in an AI lab contend, 'saliency detection', a specific image segmentation process for determining the most salient object in an image, entails making a 'really fuzzy decision' (Jatón, 2021, p. 71) even if the rote segmentation of all objects in the first place is rather less critical. Yet, in segmentation work in the domain of autonomous driving, even this initial annotation/segmentation work is important: the difference between the hood of a vehicle being delineated or not, or the back wheel of a bicycle being fully traced or not. Here, each applied domain—as well as the algorithmic process itself—starts to weigh upon the task-at-hand.

This is one of the many ways in which 'ontological politics' (Mol, 1999; Law, 2002) enters into the world of machine vision and autonomous driving: how the world is divided up, and what these divisions are subsequently referred to as. It is, in other words, a process of 'ground truthing' by which the true ground, that is the foundational basis on which ML modelling proceeds, is established.

Ground Truth, Tasks, and Metrics

Ground-truthing is a cartographic term, deriving from the process of verifying information at a specific location. Ground-truthing is thus a 'localized' process, by which human bodies and eyes can be used to 'calibrate' remotely captured and collected data (Gil-Fournier & Parikka, 2021, p. 2). By way of example, one might collect vegetation data using aerial lidar, providing the operators with a large volume of land use data otherwise impossible to collect. In order to validate such data or provide a greater contextual understanding of the remotely collected data, the operators might head to a particular site captured in the (remote) data.

Here, ground observation is a way of providing a direct, empirical account of the phenomenon captured. It is a term, and a process, commonly used by Geographical Information Science (GIS) and remote sensing practitioners. Much work in critical GIS has focused on the epistemological differences between forms of localized ground-truthing and remote positivist methods (Robbins, 2003).

That the term and practice of ground-truthing is also present and active in ML communities is interesting. Here, with no ‘outside’ world to be captured, ground-truthing in machine learning takes the form of a completely computational relation. Ground-truthing, in this sense, does not involve practitioners making their way to the site of the underlying data in order to visually, directly observe it (as remote sensors might do), but of statistically relating modelling work in relation to the so-called ‘ground truth’ of the training dataset, established as the foundation on which ML work can be built. Once rendered as an input for model training, there is little interrogation of training data itself, be it source or subsequent shape. But without the training dataset as a ground truth, there is no way in which a specific ML model or method can be judged as accurate or, indeed, successful as a method. The training dataset as a nominally certified stable, and accurate, ground offers that possibility.

Thus, a divergent critique of the politics of ground-truthing emerges between cartographic and AI settings. In critical GIS, ground-truthing is often (erroneously) valorised for its localized ‘unmediated’ epistemology (Gil-Fournier & Parikka, 2021, p. 1255), in contrast to the distant, technological eyes of remote sensing, lacking context, removed from place. Yet in ML contexts, ground-truthing contains precious little of this epistemological dimension, rendered instead as an operational necessity. Whilst ground-truthing is also fully operationalized in remote sensing processes as well, it retains this critical dimension to those in and around the discipline itself.

This is perhaps where the work of Thylstrup (2022) and others is instructive: an attempt to offer a critical perspective on training datasets at the point in which they enter the AI system ostensibly from the outside world.

As Jatón (2021) considers, ground-truthing is dependent upon two parts: a training (data)set and an evaluation (data)set. The training set is

used to train the algorithm to perform the task in hand (e.g. categorizing road users), whilst the evaluation set is used to judge its accuracy in doing so. As a result, neither dataset should contain datapoints present in the other, which would otherwise present a false picture of the performance of the ML model: fully overlapping training and evaluation sets would unsurprisingly yield a fully 100% accuracy rate, as the model would be evaluated using the same datapoints (images, etc.) as that which it was trained on. If released into the wild to be used on entirely new scenes and scenarios, there would be no sense whatsoever of its capabilities. Indeed, without further refinement it would likely perform poorly, subject to ‘overfitting’—well-attuned only to its training data.

As each domain of object-recognition is somewhat unique, with specific kinds of objects needing to be classified, such ‘high-level detection algorithms’, as Jatón (2021, p. 55) refers to them, are necessarily ‘task-specific’ (Jatón, 2021, p. 55) or task-orientated. Here, performance metrics or otherwise ‘precision and recall metrics’ (Jatón, 2021, p. 55) are critically important for determining the accuracy, and hypothetical usability or viability, or particular methods. Hence also the continued thirst for large, well-annotated training datasets in each particular domain: high-level detection algorithms that work in one domain (e.g. medical imaging) have little hope of working well in another (driving).

Locating Critique

What is perhaps different about training datasets compiled for autonomous driving purposes is that certain critiques of the way training data is collected, annotations performed, and assumptions are made, do not necessarily hold. To begin with, training datasets compiled for autonomous driving do not have a primary interest in recording faces, and therefore, little interest in tracking things like emotions or concern the array of racial, ethnic, national, professional, or behavioural categories that Crawford and Paglen discuss in reference to the ‘canonical training set’ (Crawford & Paglen, 2019, p. n.p.) ImageNet.

These training datasets are connected, however. Firstly, training datasets designed specifically for autonomous driving are indeed interested in

object classification, as discussed before. Secondly, that training datasets designed for autonomous driving applications might be used in concert with neural networks ‘pretrained’ on non-domain specific training datasets such as ImageNet. ML model diagrams contained within the technical papers submitted to autonomous driving machine vision challenges routinely show methods using pretrained neural networks on novel training data (Hind et al., 2024).

What is a critical aspect to consider, then, is that motion planning training datasets (those comprising annotated data of vehicle trajectories), offer a calculation and categorization of certain fixed properties of an individual road user/object (a car, a sign), but under certain categories (e.g. vehicle), their capacities as a *potentially moving object*. Elaine Herzberg was killed by a developmental autonomous Uber ATG vehicle partly on the evidence of her *not* being classified as a particular kind of moving road user (cyclist, rather than pedestrian), meaning that her actualized capacity as a pedestrian-walking-with-a-bicycle was neither captured nor classified, nor were her movements predicted or acted upon. Herzberg was therefore not killed because she was not ‘correctly’ classified as a pedestrian, but as a pedestrian with the actualized capacity to move across the road in a particular manner, in a particular location, at a particular speed, and in combination with other objects.

Thus, in the domain of autonomous driving, critique must be located at two moments in the compilation of training data: firstly, in the constitution of perception training data, and secondly, in the composition of motion planning training data. Whilst training datasets designed for the development of autonomous vehicles are often referred to in the singular (e.g. Waymo Open Dataset), they ordinarily contain two types, meant for the recognition of objects as well as the tracking of objects. These two practices are intertwined, without which an autonomous vehicle could not function. In this, one might directly counter the observation that ‘the material force of categories appears always and instantly’ (Bowker & Leigh Star, 1999, p. 3) and suppose almost the opposite: that the material force of the categories constructed by training datasets appear only sometimes and with delay.

The associated ML modelling performed for each is necessarily different too, as a result of the different calculative tasks at hand. Thus, whilst

the locus of responsibility can be shifted towards motion planning errors in the case of Herzberg, this should not be at the expense of removing all responsibility from the categorization/recognition locus entirely. Indeed, that following the pipeline metaphor, that without the erroneous, egregious, misclassification of Herzberg as simply ‘other’ rather than a pedestrian, the system would not have been capable of further miscalculating her trajectory resulting in her death.

KITTI Vision Benchmark Suite Versus Waymo Open Dataset

The chapter compares these two training datasets, the KITTI Vision Benchmark Suite and Waymo Open Dataset, for multiple reasons.

Firstly, the KITTI dataset was released prior to, laid the groundwork for, and arguably stimulated the recent commercial interest in autonomous driving. The KITTI Vision Benchmark Suite, as the name suggests, was central to what I call here the ‘benchmark era’ of autonomous driving. The Waymo Open Dataset, by comparison, was released 10 years after the founding of Google’s autonomous vehicle project (in 2009), three years since the founding of Waymo itself (in 2016), and one year after the fatal Uber ATG crash in Tempe, Arizona (2018) (Waymo, 2019). The KITTI dataset can be understood as the foundation for the early commercial development of autonomous driving, with the Waymo datasets understood as *a* (rather than *the*) basis for a ‘platformized’ version of autonomous driving (Hind & Gekker, 2024, Hind et al., 2022), built on the mobilization of open data for commercial development, and the use of external, competitive labour to achieve technical gains on ML tasks related to ‘perception’ (machine vision, object-recognition, semantic labelling, 2D, 3D) and motion (planning, forecasting) (Hind et al., 2024). In short, the Waymo Open Dataset can be considered the foundation for the ‘incremental era’ of autonomous driving, driven by these platformization attempts.

The KITTI dataset was compiled by researchers funded by the Karlsruhe Institute of Technology (KIT), Germany, and the Toyota

Technological Institute at Chicago (TTIC), USA, two renowned academic computer science institutions. In contrast, the Waymo datasets contain data collected by Waymo's own developmental autonomous vehicles, compiled by an in-house team.

There is also a subtle difference in the name of each project. Rather than a dataset, *per se*, KITTI is described as a 'benchmark suite'. Benchmarks are a way of comparing the performance of different ML methods, by establishing a common dataset each method can use to determine comparative performance. If different methods are trained using different datasets, comparing performance is difficult. Thus, establishing the KITTI Vision Benchmark Suite was driven by the apparent dearth of 'visual recognition systems' (Geiger et al., 2012, p. 3354) being used in robotics research, attributed to the 'lack of demanding benchmarks that mimic' (Geiger et al., 2012, p. 3354) real-world scenarios. In contrast to existing datasets derived from laboratory settings, KITTI offered a set of benchmarks in real-world, on-road settings, arguably for the first time in autonomous driving research.

The technical set-up and platform for each project was also marginally different. For the KITTI team, this involved equipping a single host vehicle—a mid-2000s Volkswagen Passat B6—with ten different sensors, including a single 'laser scanner' or lidar (Geiger et al., 2023b). Waymo, on the other hand, were able to draw on a fleet of vehicles, including a hybrid Chrysler Pacifica minivan which was retired in 2023 having been used since 2017 (Korosec, 2023) and electric Jaguar I-PACE models launched in 2018 (Waymo, 2018), both equipped with five lidar sensors and five high-resolution cameras (Sun et al., 2020). Whilst Waymo attached sensors to the front, sides, and rear of the vehicle, the KITTI team placed all ten sensors on the roof of their Volkswagen Passat. Although Waymo had previously used off-the-shelf Velodyne lidar sensors, the same HDL-64E model as used by the KITTI team (Amadeo, 2017), it switched to developing in-house sensors around 2017 (Ross, 2019).

Ground-Truthing the Domain

The KITTI dataset is derived from data captured in the city of Karlsruhe, in the south-west of Germany. Whilst described as a ‘mid-size city’ by the team behind KITTI (Geiger et al., 2023a, p. n.p.), its population of 308,000 can be considered small by both US and Chinese standards, two countries now driving commercial autonomous vehicle work. In comparison, the original release of the Waymo Open Dataset covered three US cities/urban areas: San Francisco (California), Phoenix (Arizona), and Mountain View (Google/Alphabet HQ) (Sun et al., 2020). After updating the datasets in 2021 and 2022, the data now additionally covers Los Angeles (California), Detroit (Michigan), and Seattle (Washington). All data comprising the public release of Waymo Open Datasets is derived, therefore, from US cities/predominately urban areas.

KITTI was also, arguably, the first dataset used for autonomous driving research that derived source data from a mix of lidar sensors and cameras. As Geiger et al. noted at the time, ‘visual sensors are rarely exploited in robotics applications: Autonomous driving systems rely mostly on GPS, laser range finders, radar as well as very accurate maps of the environment’ (Geiger et al., 2012, p. 3354). Thus, KITTI ushered in a new era of autonomous driving research less dependent on cartographic technologies and more dependent on sensor technologies, where data is collected from an assemblage of sensor and camera systems for the first time in an on-road, rather than off-road, environment. Interestingly, this was five years after the same KIT team had finished last in the 2007 DARPA Urban Challenge, with their vehicle AnnieWAY. The data for the KITTI dataset was collected by the same AnnieWAY vehicle, their Volkswagen Passat B6.

Particular individuals have played a significant, recurring, role in the development of autonomous driving. Raquel Urtasun is one such individual. A co-author of the technical paper explaining the KITTI dataset in 2012 (Geiger et al., 2012), and previous TTIC employee, Urtasun was the former Chief Scientist and Head of R&D at Uber ATG (2017–2021), having launched Waabi in 2021, as Chap. 2 details. The technical paper introducing the KITTI dataset (Geiger et al., 2012) is also arguably the

most cited in the autonomous driving domain, having been referenced 13,130 times.¹ It is also, by some margin, each of the co-author's (Andreas Geiger, Philip Lenz, Raquel Urtasun) most cited articles. In other words, it is a foundational technical paper describing the foundational machine vision dataset in the autonomous driving community.

There is also a substantial difference in the comparative size of the amassed datasets. The KITTI dataset consisted, at the time of launch in 2012, of 389 stereo and optical flow image pairs, 22 3D stereo image sequences totalling 39.2 km, and over 200,000 3D object annotations in what they describe as 'cluttered scenarios' (Geiger et al., 2012, p. 1) with up to 15 vehicles and 30 pedestrians in any one image. According to the KITTI team's own comparisons, no other dataset up until that point had captured more than 6.4 km of stereo sequences (the 'Málaga 2009 dataset' collected by researchers at the University of Málaga) (Blanco et al., 2009), and only one team at the Technical University of Munich (TUM) had generated more sequences (27–22), although critically, these had only been generated indoors rather than on public roads (Sturm, 2017). Thus, while the KITTI dataset was by no means the first such dataset, it established a new benchmark for autonomous vehicle training datasets. By comparison, when first launched in 2019, the Waymo Open Dataset consisted of 1150 scenes—what the KITTI team call scenarios—each totalling 20 seconds. This constituted a 50-fold increase in available 3D stereo scenes/scenarios for training. The Waymo dataset comprises 12 million 3D bounding boxes compared to 80,000 in the KITTI dataset, constituting 6.4 hours of captured data (Sun et al., 2020).

Excavating the Interesting

Most notably, Waymo claimed that the original Open Dataset was '15x more diverse than the largest camera+LIDAR dataset available' (Sun et al., 2020, p. 1) based, that is, on a particular geographical metric devised by Waymo to evaluate scenario diversity. Here 'dataset diversity' is not a term explicitly used by the KITTI team, whereas Waymo refer to urban/suburban distinctions, 'time of day diversity' (Sun et al., 2020, p. 5) and scenes 'selected from many different parts' (Sun et al., 2020,

¹ At the time of writing (December 20, 2023).

p. 5) of Phoenix (40 km² total), San Francisco, and Mountain View (36km² combined). In the conclusion to their technical paper, the Waymo authors once again state that the Open Dataset is ‘significantly larger, higher quality, more geographically diverse than any existing similar dataset’ (Sun et al., 2020 p. 8), with an unrivalled level of so-called domain diversity (Sun et al., 2020, p. 8) amongst the captured data from the three aforementioned locations. Whilst the KITTI team do in fact talk about diversity, this strictly concerns the technical work being performed: (a) in relation to selecting a subset of the training dataset for evaluation purposes and (b) in relation to the metrics used to perform such evaluation. Thus, diversity is present in four senses: sensor diversity, domain diversity, evaluation diversity, and metric diversity. All these are brought to bear on the datasets in fascinating ways.

This emphasis on diversity, diverse environments, and diverse conditions is also matched by an interest in ‘interactive situations’ and ‘interesting interactions’ (Ettinger et al., 2021, p. 1). In this, the quality of ‘interesting-ness’ is derived from the manner in which different objects or agents represented in training data interact with each other, critical ‘to develop motion planning models’ (Ettinger et al., 2021, p. 1). As Waymo researchers have suggested, ‘of particular importance are interactive situations such as merges, unprotected turns, etc., where predicting individual object motion is not sufficient’ (Ettinger et al., 2021, p. 1). Likewise, they talk of ‘mining for interesting scenarios’ (Ettinger et al., 2021, p. 3), suggesting that such scenarios are buried beneath a greater volume of scenarios deemed ‘not interesting’ at all. As Ramon Amaro suggests, with the help of Jiawei Han et al. (2011), this search for interestingness ‘reveals the fragility of machine perception as a function of human desire and expectation’ (Amaro, 2022, p. 123). Not all patterns are interesting, as Han et al. (2011) likewise contend, and certain methods might be used to help determine interestingness in ML datasets. These include objective methods such as associational rule *support* (number of connections) and associational rule *confidence* (degree of certainty of connection), each establishing thresholds regarding associations between entries in a dataset. Interesting scenarios, therefore, might be statistically determined: those where rule support or confidence values surpass a set threshold of interaction between objects or agents. Likewise, more subjective methods

might help to determine whether patterns or associations are ‘unexpected’ or ‘actionable’ (Han et al., 2011, p. 22) and hence deemed interesting.

Scott Ettinger, lead author of the Waymo motion forecasting technical paper where the notion of interesting-ness is discussed, was a member of the winning 2005 DARPA Grand Challenge team, Stanford Racing Team (Thrun et al., 2007), demonstrating the certain continuity between these otherwise distinct phases or eras of autonomous driving. The training dataset, therefore, can be understood as a ‘device’ for the ‘automated production of *interested readings of empirical reality*’ (Rieder, 2020, p. 252, authors’ emphasis), where the empirical substrate to be worked upon is considered interesting enough to constitute an empirical reality underfoot.

Interestingness can be understood as not only a variation of the theme of dataset diversity but another way of talking about ‘edge cases’ in training datasets. Edge cases are instances that are not common to a particular dataset, only likely to be expressed in very particular, extreme, or ‘edge’ situations. Whilst uncommon, edge cases assume great importance within the development of autonomous vehicles, particularly because of their ability to cause great harm if actually expressed.

Part of this quest for interestingness—to produce ‘interested readings of empirical reality’—also pertains to the interactions themselves. Whilst it may be relatively straightforward to consider what kinds of interactions occur in driving situations, *differentiating* between and *quantifying* types of interactions is altogether more difficult. Here, the same Waymo researchers draw on work by colleagues on ‘conditional behaviour prediction’ (CBP) (Tolstaya et al., 2021) to determine the ‘degree of influence’ (Ettinger et al., 2021, p. 7) one agent has on another. What this ultimately means is that interestingness is defined through a quantification of interaction itself. Interestingness is complexity, where complexity is the extent to which multiple agents are said (quantified) to be interacting with each other.

Neither training dataset, in some cases quite literally, is a static proposition. Whilst the KITTI dataset was launched in 2012, now over 12 years old, new benchmarks comprising of 200 training scenes and 200 test scenes were released in 2015 (Geiger et al., 2023c). One key difference was that in contrast to the (static) images pairs from 2012, these constituted ‘highly dynamic scenes’ (Menze & Geiger, 2015, p. 1), consisting of moving images. Whilst drawn from the same raw KITTI dataset from 2012, these

new benchmarks constituted an important development, allowing other researchers to use such benchmarks to train their own motion planning models, what the authors refer to in the technical paper as ‘object scene flow’ (Menze & Geiger, 2015, p. 2). In 2015, this work was decidedly new, with the authors contending that ‘none of the existing optical flow or scene flow methods [were] able to cope with the extreme motions produced by moving objects in some of our scenes’ (Menze & Geiger, 2015, p. 7). By comparison, Waymo’s Open Dataset has been updated multiple times, ordinarily in line with each annual Open Dataset Challenge. In 2021, they published a motion dataset to compliment a perception dataset, in 2022 they added additional labels to both existing datasets, and in 2023 they announced the inclusion of more sensor data, new features to the datasets, and an entirely new ‘modular data format’ to more ‘efficiently access and explore [the] data’ (Waymo, 2023, p. n.p.).

Incremental Gains

One final aspect is the labour involved in annotating the respective datasets to allow their interesting qualities to be surfaced. Whilst many training datasets are annotated by forms of piecemeal, distributed, online labour, the canonical datasets in the domain of autonomous driving have typically been annotated by associated researchers, rather the external workers. As Cordts et al. write, on the Cityscapes Dataset:

Our 5000 fine pixel-level annotations consist of layered polygons ... and were realized in-house to guarantee the highest quality levels. Annotation and quality control required more than 1.5h on average for a single image. Annotators were asked to label the image from back[ground] to front[ground] such that no object boundary was marked more than once. Each annotation thus implicitly provides a depth ordering of objects in the scene. (Cordts et al., 2016, p. 3214)

The same can be said for the Waymo Open Dataset, compiled and annotated in-house, as well as the other autonomous driving training datasets listed in Table 3.1. What is interesting to note, therefore, is that

Table 3.1 A comparison between autonomous driving datasets

Dataset name	Funders/ organizations	Volume	Data type	Task	Collection method
Cityscapes Dataset	TU Darmstadt, Max Planck Institute for Informatics, Daimler AG, TU Dresden, German Federal Ministry for Economy and Technology (BMW)	5000 (images fine) 20,000 (images, coarse), ~1000 (video, fine)	Camera, video	Semantic understanding	Single vehicle
KITTI Vision Benchmark Suite	Karlsruhe Institute of Technology (KIT), Toyota Technological Institute at Chicago (TTIC)	194 training scenes, 195 test scenes (2012), 200 training scenes, 200 test scenes (2015)	Stereo camera	Perception (object- recognition)	Single 'AnnieWAY' Volkswagen Passat B6
Waymo Open Dataset	Waymo (Google/ Alphabet)	1000 driving segments, 12 million 3D labels, 1.2 million 2D labels (2019)	Lidar, camera	Perception, motion planning	Modified fleet (Jaguar I-Pace, Chrysler Pacifica)
Argo Argoverse	Argo AI, Carnegie Mellon University (CMU), Georgia Institute of Technology (GT), Ford	300,000 5-second scenarios	Lidar, 360°, stereo camera	3D tracking, forecasting	Modified fleet (Ford Fusion Hybrid)

Location(s)	Release Date	Updates	Autonomous driving project application	Technical paper	Technical paper citations
50 cities (Germany, Switzerland)	20/02/2016	2020	Various	Cordts et al. (2016) <i>The Cityscapes Dataset for semantic urban scene understanding</i>	9266
Karlsruhe (Germany)	20/03/2012	2015	Various	Geiger et al. (2012) <i>Are we ready for autonomous driving? The KITTI Vision Benchmark Suite</i>	13130
San Francisco, Phoenix, Mountain View, Kirkland (original, 2019), Los Angeles, Detroit, Seattle (added, 2021)	21/08/2019	2021, 2022, 2023	Waymo	Sun et al. (2020) <i>Scalability in perception for autonomous driving: Waymo Open Dataset</i>	2097
Pittsburgh, Miami (original, 2019), Austin, Washington DC, Palo Alto, Detroit (added, 2021)	19/06/2019	2021	Argo AI	Chang et al. (2019) <i>Argoverse: 3D tracking and forecasting with rich maps</i>	1130

the work underpinning these datasets is not derived from platform labour, or piecemeal labour, in the form that has come to be associated with training datasets—even if the ML model training work has become externalized, subject to platformized relations (Hind et al., 2024).

This annotation work is therefore not the kind of distributed, online ‘micro-work’ others have uncovered in the automotive industry, connected to training voice assistants (Tubaro & Casilli, 2019), even though it can be considered a derivative form of such ‘computer-supported on-demand low-valued work’ (Jaton, 2021, p. 66). Indeed, it is perhaps best thought of the other way around: that in-house annotation work can reasonably be understood to be the precursor to, and foundation for, the subsequent distribution and platformization of AI annotation work. As AI moved out of the laboratory and into the workplace, so annotation work found a new home elsewhere: predominantly, although not exclusively, in the homes of women in the global south, Latin America, and countries with enduring postcolonial legacies (Tubaro et al., 2022; Viana Braz et al., 2023).

Yet, due to the stated importance of the annotation work at hand, ‘to guarantee the highest quality levels’, the work undertaken to annotate images for autonomous driving is typically carried out by those with a pre-existing knowledge of the subject area and an ongoing interest in ensuring the quality of annotations, beyond piecemeal payment. In other words, that the annotators are, in principle, also the machine learners, and rather than being alienated from one side of the equation, annotators know precisely what the later, realizable value of their annotation work is, and what their annotations are going towards. As often stated, micro-workers commonly do not know why they are extracting information from image inputs, nor what their annotation work is necessarily going towards. Whilst sometimes meaning can be deduced by workers, some annotation work is plainly mystifying (Tubaro et al., 2020). In the case of KITTI and Waymo Open Datasets, this ‘AI preparation’ work, as Tubaro et al. (2020, p. 5) understand it, offers a vastly different experience for those annotating. It is also, as one might expect, work being undertaken by those paid vastly higher amounts, essentially on computer science PhD stipends or scholarships, funded in part by big tech or automotive firms.

But what is the value or need for surfacing the interesting qualities of the respective training datasets? Increasingly, as the Waymo Open Dataset demonstrates, it is to deliver incremental gains in the performance of ML models. Aided by their Waymo Open Dataset Challenges (2020–2023), their training datasets have offered the possibility to deliver small, but no less significant, performance gains on object-recognition related tasks (Hind et al., 2024). Put to work, Waymo’s Open Datasets provide the empirical foundation on which such model training is done, structuring competition between rival model development teams. Here, incremental gains are delivered in respect to certain object-recognition metrics such as ‘average precision’ (AP), used to assess the accuracy of the ML-based object-recognition techniques developed. Relatively small gains in the accuracy of winning methods in the annual Open Dataset Challenges were recorded, despite a huge number of varied entrants (Hind et al., 2024). As a team of computer scientists considered following the end of a precursor competition, such challenges pose the risk of ‘reduc[ing] the diversity of methods within the community’ as truly novel methods are jettisoned ‘before they have the chance to mature’ (Everingham et al., 2015, p. 133). Concentration in both available training datasets and ML models trained on them ultimately entrenches incrementalist approaches—a hallmark of the current era of autonomous vehicle development.

Conclusion

This chapter has narrated the development of two autonomous vehicle training datasets, the KITTI Vision Benchmark Suite and Waymo Open Dataset. Whilst just two amongst many others, these training datasets represent key milestones in the development of autonomous vehicles. The KITTI Vision Benchmark Suite laid the groundwork for the testing of object-recognition and motion planning systems from the early 2010s onwards. Waymo’s public release of their Open Dataset in 2019 can be seen as a landmark in the platformization of autonomous driving (Hind et al., 2022).

Across these datasets and timeframes, however, are a shared interest in the interesting. Training dataset diversity, as this chapter has argued, is a critical, ongoing concern for machine vision researchers working in the domain of autonomous driving. How this diversity is expressed—whilst sharing generalities across domains—is specific to the particularities of designing object-recognition and motion planning systems for navigating driving environments. In this, the search for dataset diversity is expressed through references to diverse situations, scenes, and scenarios with the desire for so-called interesting interactions to be represented within such datasets. Everything from the location of the original capture of training data—be it Karlsruhe, Germany, or Phoenix, Arizona—to the volume of recorded segments of video footage has a role in ensuring maximum diversity in the datasets being composed.

Invariably, however, as the ‘mining’ metaphors employed by some researchers suggest, the search for interesting interactions requires considerable work. In other words, driving manoeuvres deemed valuable by engineers—from merges to unprotected turns—or driving encounters between different road users (pedestrians and cars, cyclists and cars) do not always present themselves within the data but must be expressed and surfaced. Lacking a workable volume of these interactions in the datasets being used to train object-recognition and motion planning ML models, greater risks in real-world scenarios (and real-world interactions) are inevitably generated down the line.

Thus, different kinds and levels of training data work are important, from the fastidious work of annotating training data to the beneficial quantification of interactivity. In each case, ‘mining for interesting scenarios’ (Ettinger et al., 2021, p. 3) always requires a substantial pool of able and willing ‘miners’, those equipped to implement routine processes as well as develop novel methods and strategies for designing autonomous vehicle systems.

References

- Amadeo, R. (2017, January 10). Google's Waymo invests in LIDAR technology, cuts costs by 90 percent. *Ars Technica*. <https://arstechnica.com/cars/2017/01/google-s-waymo-invests-in-lidar-technology-cuts-costs-by-90-percent/>
- Amaro, R. (2022). *The black technical object: On machine learning and the aspiration of black being*. Sternberg Press.
- Bechmann, A., & Bowker, G. C. (2019). Unsupervised by any other name: Hidden layers of knowledge production in artificial intelligence on social media. *Big Data & Society*, 6(1), 1–11. <https://doi.org/10.1177/2053951718819569>
- Blanco, J. L., Moreno-Dueñas, F. A., & González, J. (2009). The Málaga 2009 robotic dataset collection. *MAPIR Lab*. Retrieved October 1, 2023, from https://www.mrpt.org/downloads/dataset2009/README_campus_2L.html
- Bowker, G. C., & Leigh Star, S. (1999). *Sorting things out: Classification and its consequences*. MIT Press.
- Chang, M.-F., Lambert, J., Sangkloy, P., Singh, J., Bæk, S., Hartnett, A., Wang, D., Carr, P., Lucey, S., Ramanan, D., & Hay, J. (2019). Argoverse: 3D tracking and forecasting with rich maps. Preprint. <https://doi.org/10.48550/arXiv.1911.02620>
- Cityscapes. (2022). Dataset overview. *Cityscapes Dataset*. Retrieved October 1, 2023, from <https://www.cityscapes-dataset.com/dataset-overview/>
- Cordts, M., Omran, M., Ramos, S., Rehfeld, T., Enzweiler, M., Benenson, R., Franke, U., Roth, S., & Schiele, B. (2016). The Cityscapes Dataset for semantic urban scene understanding. *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016, 3213–3223. <https://doi.org/10.1109/CVPR.2016.350>
- Crawford, K., & Paglen, T. (2019, September 19). Excavating AI: The politics of training datasets for machine learning. *Excavating AI*. <https://excavating.ai/>
- Ettinger, S., Cheng, S., Caine, B., Liu, C., Zhao, H., Pradhan, S., Chai, Y., Sapp, B., Qi, C., Zhou, Y., Yang, Z., Chouard, A., Sun, P., Ngiam, J., Vasudevan, V., McCauley, A., Shlens, J., & Anguelov, D. (2021). Large scale interactive motion forecasting for autonomous driving: The Waymo Open Motion Dataset. Preprint. <https://doi.org/10.48550/arXiv.2104.10133>
- Everingham, M., Ali Eslami, S. M., Van Gool, L., Williams, C. K. I., Winn, J., & Zisserman, A. (2015). The PASCAL Visual Object Classes Challenge: A retrospective. *International Journal of Computer Vision*, 111, 98–136. <https://doi.org/10.1007/s11263-014-0733-5>

- Geiger, A., Lenz, P., Stiller, C., & Urtasun, R. (2023a). Welcome to the KITTI Vision Benchmark Suite! *The KITTI Vision Benchmark Suite*. Retrieved October 1, 2023, from <http://www.cvlibs.net/datasets/kitti/>
- Geiger, A., Lenz, P., Stiller, C., & Urtasun, R. (2023b). Sensor setup. *The KITTI Benchmark Suite*. Retrieved October 1, 2023, from <https://www.cvlibs.net/datasets/kitti/setup>
- Geiger, A., Lenz, P., Stiller, C., & Urtasun, R. (2023c). Stereo evaluation 2015. *The KITTI Benchmark Suite*. Retrieved October 1, 2023, from https://www.cvlibs.net/datasets/kitti/eval_scene_flow.php?benchmark=stereo
- Geiger, A., Lenz, P., & Urtasun, R. (2012). Are we ready for autonomous driving? The KITTI Vision Benchmark Suite. *2012 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 3354–3361. <https://doi.org/10.1109/CVPR.2012.6248074>
- Gil-Fournier, A., & Parikka, J. (2021). Ground truth to fake geographies: Machine vision and learning in visual practices. *AI & Society*, 36, 1253–1262. <https://doi.org/10.1007/s00146-020-01062-3>
- Han, J., Kamber, M., & Pei, J. (2011). *Data mining: Concepts and techniques*. Morgan Kaufmann.
- Hind, S., Kanderske, M., & van der Vlist, F. (2022). Making the car ‘platform ready’: How big tech is driving the platformization of automobility. *Social Media + Society*, 8(2), 1–13. <https://doi.org/10.1177/20563051221098697>
- Hind, S., van der Vlist, F., & Kanderske, M. (2024). Challenges as catalysts: How Waymo’s Open Dataset Challenges shape AI Development. *AI & Society*, 0(0), 1–17. <https://doi.org/10.1007/s00146-024-01927-x>
- Hind, S., & Gekker, A. (2024). Automotive parasitism: Examining Mobileye’s ‘car-agnostic’ platformisation. *New Media & Society*, 26(7), 3707–3727. <https://doi.org/10.1177/14614448221104209>
- Jaton, F. (2021). *The constitution of algorithms: Ground-truthing, programming, formulating*. MIT Press. <https://doi.org/10.7551/mitpress/12517.001.0001>
- Korosec, K. (2023, March 30). Waymo retires its self-driving Chrysler Pacifica minivan. *Techcrunch*. Retrieved October 1, 2023, from <https://techcrunch.com/2023/03/30/waymo-retires-its-self-driving-chrysler-pacifica-minivan/>
- Law, J. (2002). *Aircraft stories: Decentering the object in technoscience*. Duke University Press.
- Mackenzie, A. (2017). *Machine learners: Archaeology of a data practice*. MIT Press.
- Menze, M., & Geiger, A. (2015). Object scene flow for autonomous vehicles. *CVPR 2015*. Retrieved October 1, 2023, from <https://www.cvlibs.net/publications/Menze2015CVPR.pdf>

- Mol, A. (1999). Ontological politics: A word and some questions. *The Sociological Review*, 47(1), 74–89. <https://doi.org/10.1111/j.1467-954X.1999.tb03483.x>
- Rieder, B. (2020). *Engines of order: A mechanology of algorithmic techniques*. Amsterdam University Press. <https://doi.org/10.2307/j.ctv12sdvfl>
- Robbins, P. (2003). Beyond ground truth: GIS and the environmental knowledge of herders, professional foresters, and other traditional communities. *Human Ecology*, 31(2), 233–253. <https://www.jstor.org/stable/4603469>
- Ross, P. E. (2019, March 7). Waymo will sell its short-range lidar, challenging industry leader Velodyne. *IEEE Spectrum*. <https://spectrum.ieee.org/waymo-will-sell-its-shortrange-lidar-challenging-industry-leader-velodyne>
- Srnicek, N. (2022). Data, computer, labor. In M. Graham, & F. Ferrari (Eds.), *Digital work in the planetary market* (pp. 241–261). MIT Press. <https://doi.org/10.7551/mitpress/13835.003.0019>
- Sturm, J. (2017). RGB-D SLAM dataset and benchmark. *Computer Vision Group, TUM School of Computation, Information and Technology*. Retrieved October 1, 2023, from <https://cvg.cit.tum.de/data/datasets/rgbd-dataset>
- Sun, P., Kretschmar, H., Dotiwalla, X., Chouard, A., Patnaik, V., Tsui, P., ..., & Anguelov, D. (2020). Scalability in perception for autonomous driving: Waymo Open Dataset. Preprint. <https://doi.org/10.48550/arXiv.1912.04838>
- Thrun, S., Montemerlo, M., Dahlkamp, H., Stavens, D., Aron, A., Diebel, J., et al. (2007). Stanley: The robot that won the DARPA Grand Challenge. In M. Buehler, K. Iagnemma, & S. Singh (Eds.), *The 2005 DARPA grand challenge: The great robot race* (pp. 1–43). Springer.
- Thylstrup, N. B. (2022). The ethics and politics of data sets in the age of machine learning: Deleting traces and encountering remains. *Media, Culture & Society*, 44(4), 655–671. <https://doi.org/10.1177/01634437211060226>
- Tolstaya, E., Mahjourian, R., Downey, C., Varadarajan, B., Sapp, B., & Anguelov, D. (2021). Identifying driver interactions via conditional behaviour prediction. Preprint. <https://doi.org/10.48550/arXiv.2104.09959>
- Tubaro, P., & Casilli, A. A. (2019). Micro-work, artificial intelligence and the automotive industry. *Journal of Industrial and Business Economics*, 46(3), 333–345. <https://doi.org/10.1007/s40812-019-00121-1>
- Tubaro, P., Casilli, A. A., & Coville, M. (2020). The trainer, the verifier, the imitator: Three ways in which human platform workers support artificial intelligence. *Big Data & Society*, 7(1), 1–12. <https://doi.org/10.1177/2053951720919776>

- Tubaro, P., Coville, M., Le Ludec, C., & Casilli, A. A. (2022). Hidden inequalities: The gendered labour of women on micro-tasking platforms. *Internet Policy Review*, 11(1), 1–26. <https://doi.org/10.14763/2022.1.1623>
- Viana Braz, M., Tubaro, P., & Casilli, A. A. (2023). Microwork in Brazil: Who are the workers behind artificial intelligence? *Research Report DiPLab & LATRAPs*. Retrieved October 1, 2023, from <https://diplab.eu/?p=2833>
- Waymo. (2018, March 27). Meet our newest self-driving vehicle: The all-electric Jaguar I-PACE. *Waymo*. <https://waymo.com/blog/2019/08/meet-our-newest-self-driving-vehicle.html>
- Waymo. (2019, August 19). Waymo Open Dataset: Sharing our self-driving data for research. *Waymo*. <https://blog.waymo.com/2019/08/waymo-open-dataset-sharing-our-self.html>
- Waymo. (2023, March 16). Driving research forward: The Waymo Open Dataset updates and 2023 challenges. *Waymo*. <https://waymo.com/blog/2023/03/driving-research-forward-waymo-open.html>



4

Sensing Decisions: Perceiving, Classifying, Finessing

As the previous chapter considered, autonomous vehicles are reliant upon training datasets to provide a ‘ground truth’ for their eventual operations. Such training datasets themselves—like the KITTI Vision Benchmark Suite—have necessarily been compiled with the help of vehicles equipped with an array of sensors able to capture the training data itself. A considerable amount of labour is subsequently involved in calibrating, annotating, and making available such training data to those that want to use it. Their broad aim is to train image-based machine learning (ML) models to recognize objects—a process that underpins simulation work to follow.

This chapter will consider the ‘sensor strategies’ devised by these machine vision researchers to improve the performance of computational tasks concerning perception and classification. Returning to the introduction, the chapter explores how this ‘sensor work’ is not only integral to the future operation of autonomous vehicles but critical to ensuring the uninterrupted flow of sensor data from module to module, site to site. In this, sensing decisions enable the fundamental ‘inter-operability’ of autonomous vehicles, with the ‘finessing’ (Fisch, 2018, p. 29) of specific machine vision strategies, techniques, and methods part-and-parcel of delivering autonomous driving.

Sensor data—whether in the form of radar returns, lidar point clouds, camera images, or other externally captured, internally transmitted, data—is integral to the decision-making capabilities of autonomous vehicles. Data generated and processed by an oft-dizzying array of sensors nominally allows any such vehicle to perceive the environment around it, in which sensor data is generated by different sensor *types* (lidar, radar, camera, etc.) and sensor *units* (Velodyne HDL-64E, etc.), in different sensor data *formats*. Some sensing systems (such as radar) allow vehicles to perceive more distant phenomena, whilst others (such as lidar) facilitate depth perception or enable the recognition of specific objects (such as cameras). This is what I have called the *composite* and *distributed* nature of sensing (Hind, 2022).

I begin by drawing on literature within media studies on the production of ‘operative images’ (Farocki, 2004; Hoel, 2018; Distelmayer, 2018) and ‘operational data’ (Walker Rettberg, 2020), categories distinct from so-called representational images or data through their utilization within, and by, technological systems. Following the above, however, I argue that sensor data used in the training of ML models should instead be conceived as *interoperable* and integral to the *interoperation* of autonomous vehicles—rather than simply being understood as ‘operative’ or operational, per se. I discuss what is meant by interoperability, with respect to autonomous vehicles, and outline the levels or stages of such interoperation, necessarily involving interoperability at both technical and epistemic levels (Wilmott, 2016, 2020).

To evidence this I draw on participant observation of two machine vision events: a virtual summit called ‘Machines Can See’ held in June 2020, and a ‘Workshop on Autonomous Driving’ (WAD) hosted in June 2021, as part of the annual Computer Vision and Pattern Recognition (CVPR) conference, a premier academic conference in the world of machine vision. At both events cutting-edge object-recognition work was presented, including techniques such as ‘3D object detection’ and ‘streaming perception’ that I understand as enabling greater interoperability within autonomous vehicles.

As a way to understand these techniques, I draw on the work of anthropologist Michael Fisch (2018), suggesting that machine vision experts ‘finesse’ interoperability; employing and refining certain skills learnt in their time as researchers in the field.

Operational: Images, Data, Clouds

Sensor data generated by autonomous vehicles is formatted in particular ways. Whether as radar returns, point clouds, or video frames, this sensor data can be said to become ‘operational’ (Mackenzie & Munster, 2019; Walker Rettberg, 2020), generating ‘operative images’ (Farocki, 2004; Hoel, 2018; Distelmayer, 2018) derived from such sensor data. In other words, they are enrolled into the machinic capacities of autonomous vehicles, such that without them, they cease to operate—effectively, efficiently, or safely: offering what Luciana Parisi (2020, p. 3) refers to as a ‘cybernetic model of steering conduct’. In this section, then, I discuss the significance of ‘operations’ and how sensor data is made operational in autonomous vehicles.

As Jill Walker Rettberg (2020) suggests, ‘operational data’ can be categorized in two ways. Firstly, data can become operational as it is depersonalized, anonymized, and aggregated. For example, in how tracking data generated by Strava users is packaged and sold to cities through a programme called ‘Strava Metro’ (Walker Rettberg, 2020, p. 8). Here, whilst such data is still visible to individual users through their own dashboards in a more ‘representational’ form, affecting and modulating the future possible activities of Strava users, ‘the representations are not the end goal’ (Walker Rettberg, 2020, p. 8). Instead, this data is made use of operationally, say, ‘as city planners use data from Strava users to redesign the city to encourage or alter the usage patterns they observe’ (Walker Rettberg, 2020, p. 8). In this, the data becomes somewhat detached from this original, representational use; circulating in other ways as it is combined and connected with other sources. Most notably, this data becomes operational in how it is used by other kinds of human actors such as the city planner or the transport official, who might come to rely on such data to make strategic decisions regarding city zoning or the installation of bike lanes.

Secondly, data can become operational when it is principally used by machines. For example, in how tracking data generated by Strava users ‘can be used as data sets for systems [to] calculate the best route for a self-driving car’ (Walker Rettberg, 2020, p. 9), dependent perhaps on

‘changing traffic patterns at different times of the day’ (Walker Rettberg, 2020, p. 9). What is different here is that such operational data is ‘algorithmically processed ... with little human involvement’, with ‘no need for human-readable representations’ (Walker Rettberg, 2020, p. 9). Indeed, that such data assumes a machine-readable form and is used for various, possible ML tasks, acting as training data on which to train algorithms to act in particular ways when faced with similar data ‘in the wild’. Whilst I want to argue here that even operational data in this second sense retains a considerable element of human involvement, either through forms of ‘supervision’ (Hind, 2019) or more fully through a ‘finessing’ of interoperability, Walker Rettberg’s categorization helps to pull apart the different ways in which data can be seen to be made operational.

Operative images are set in contrast to representational images—landscape paintings, still life, or even selfies, perhaps. These images, created in various media—paint, pencil, pixel—are not in any sense ‘working images ... tied to specialized tasks’ (Hoel, 2018, p. 12). They are principally aesthetic objects, reflective of places, things, or people. This is not to diminish or erase the obvious work involved in creating such images. But in comparison, operative images are integral to the operation of particular processes within the machine age in ways that representational images are not. Indeed, that such images are created by, and principally intended for, machines—whether in the form of drone strikes, facial recognition processes at the border, or entry to a sporting or music event. Here, the aesthetic dimension of the image recedes (although does not entirely disappear), to be usurped by a strictly instrumental requirement, to which the image’s operational dimension, and meaning (Bunz, 2019), is critical. Instead, through processes of deep learning, algorithmic systems ‘pick apart any given image into component shapes, gradients, luminosities, and corners’ (Parisi, 2020, p. 5), generating a different kind of demonstrably non-representational value. Whilst the human is never completely taken out of the loop, with machinic processes still requiring human involvement, the ‘involvement’ entailed looks distinctly different, as types of machine supervision proliferate.

However, the argument I want to make is not that all sensor data generate operative images. Sensor data is neither an operative image in itself

nor necessarily an ‘image’ at all—requiring substantial processing and formatting to be turned into anything resembling such. With lidar data, for instance, ‘point clouds’ are generated, but only through the help of secondary software. Even then, it is not at all clear if such point clouds satisfy all criteria as an image but instead may be considered as ‘operational clouds’, as Amoores (2018) or Halpern (2014) might consider them. In any case, the images produced that purport to show lidar-in-action, are dressed up, coloured, and calibrated for human eyes, as in some techniques (Hind, 2023). They are an attempt to translate point clouds into representational images, rather than an acknowledgement of their operative functionality.

To illustrate this, consider the HDL-64E, a lidar unit manufactured by Velodyne, whose founders (David and Bruce Hall) had entered the 2005 DARPA Grand Challenge, using a vehicle equipped with a prototype lidar sensor which later evolved into the HDL-64E (Velodyne, 2017, 2021). The unit itself is capable of producing ‘viewable data’ (Velodyne, 2019, p. 9) without prior configuration or calibration. Indeed, that it will ‘start scanning and producing data packets’ (Velodyne, 2019, p. 9) from the moment it is plugged in. However, in order to actually view such data, ‘point-cloud processing data viewer software’ (Velodyne, 2019, p. 9) is required. Conveniently, such software is included in the box, in the form of Velodyne’s own Digital Sensor Recorder (DSR) application, but more experienced users can elect to develop their own ‘application-specific’ (Velodyne, 2019, p. 9) viewer. In any case, Velodyne’s standard DSR application ‘reads in the packets from the sensor over Ethernet, performs the necessary calculations to determine point locations and then plots the points in 3D on [a] PC monitor’ (Velodyne, 2019, p. 9). Two critical dimensions of such lidar data, distance and intensity, are both observable through such a viewer. Thus, the point cloud is only ever a product of interoperation between a lidar unit capable of capturing and storing the sensor data, and a viewing application capable of reading and rendering the sensor data. The point cloud or ‘operational cloud’ only ever crystallizes at the point at which the unit and viewer interoperate. Neither the unit in its operation is capable of rendering a point cloud (only capturing the constituent sensor data), nor the viewer capable of capturing the sensor data (only providing the environment in which to

render the sensor data). Either acting alone is otherwise useless at creating, or actualising, such a point cloud.

As Distelmayer (2018, p. 62) contends, '[o]perative images are ... parts and thresholds of (at least) four types of mutually connected operations'. These types, he continues, involve operations between hardware and software, networked computers, computers and 'non-computer forms' (Distelmayer, 2018, p. 62), or humans and computers. For Distelmayer, then, operative images are part of what he calls 'interface operations' (Distelmayer, 2018, p. 62), connecting any one of the forms (software, hardware, networked computers, non-computer forms, humans). However, I argue here that operative images do more than work as thresholds or interfaces between other technical and non-technical objects.

Distributing Sensing, Formatting Interoperability

Instead, I want to argue that all sensor data generated by, and in, autonomous vehicles is both (a) interoperable (with the capacity to interoperate) and (b) interoperative (actually interoperating). In short, to operate is principally to interoperate. Without interoperation, operation itself cannot occur: interoperation is the *default operational state* of any autonomous vehicle. Yet, as Hoel (2018, p. 12) has suggested, 'the notion of operation is under-theorized as a media-theoretical concept, since in many cases it is simply imported from other research fields, such as computer science'.

Interoperation can be understood as a specifically media-theoretical term, able to account for the work being performed to develop the object-recognition capabilities of autonomous vehicles but also to connect a number of developmental problems—such as those discussed later—facing machine vision researchers. Rather than working with operative images 'shielded from the human eye' (Uliasz, 2020, p. 6), such practitioners are exposed to them as they flow from system to system, that is in their interoperative capacity.

An approach informed by Hoel's (Hoel, 2018) concern, then, develops a terminology around interoperability and interoperation that can ground these higher-level, abstracted, concerns. A focus on interoperation establishes the autonomous vehicle as a principle site for the interoperation of sensing technologies, algorithmic processes, *and* the movement of physical components. Starting instead from the interoperation of these entities, rather than from specific sensing technologies (such as lidar), algorithmic processes (such as object-recognition), or physical components (such as the steering wheel), removes the risk of overstating the role of any of these entities (as, say, off-the-shelf objects); foregrounding instead what they do in concert with each other. How, in other words, they might process, optimize, detect, and segment *together* rather than apart.

In this specific context then, I have defined interoperability as 'the transmission of sensor data from one (sensing) system to other, connected systems deemed necessary for subsequent decision-making processes' (Hind, 2023, p. 4) within any autonomous vehicle. Likewise that, unlike 'co-operation' interoperability does not demand that respective sensing and decision-making systems work together 'to achieve a mutual aim' (Hind, 2023, p. 4) but 'interoperate to achieve specific modular, or parallel, goals such as detecting 3D objects or processing video frames' (Hind, 2023, p. 4). Here, as outlined in the introduction, such tasks can be considered part of an operational 'pipeline' along which sensor data flows.

The notion of interoperability arguably lies behind Mackenzie and Munster's (Mackenzie & Munster, 2019) concept of 'platform seeing', where 'massive flows and iterations of images across and within devices' are facilitated by platforms and ML models, respectively, 'plat-formatted in operation' (Mackenzie & Munster, 2019, p. 9, authors' emphasis). In such cases, like contemporary smartphone cameras, 'seeing' is not performed by a single unit or component inside, capable of capturing images purely from within, as a property of the component itself. Instead, it 'is performed by a multitude of human and computational agents whose "vision" passes across and along platforms, eluding any singular coordinating position' (Mackenzie & Munster, 2019, p. 9). For something like the Apple iPhone 14 Pro, this photographic capability is made possible only through a combination of different systems ranging from

‘sensor-shift optical image stabilisation’ (OIS) to the company’s ML-driven ‘photonic engine’ (Apple, 2023).

Mackenzie and Munster’s idea of being ‘plat-formatted’ is connected to the nature of formats and the practice of formatting more generally (Volmar et al., 2020). Here the sensing capacities of the smartphone are only made possible through ‘structural or programmatic relationships between individual elements and their organizational logic’ (Volmar et al., 2020, p. 8). Whether using the term ‘format’, ‘protocol’, or otherwise, formatting requires the setting or calibration of different technical objects that, in some way, must interoperate with each other—a tape in a tape player, a CD in a CD player, a DVD in a DVD player. As Volmar et al. offer, a format describes a ‘coherent pattern of order and composition – a standardized template for the organization of space, time or information according to some rhythmical, structural, aesthetic or volumetric rules’ (Volmar et al., 2020, p. 14). Straightforwardly then, the smartphone camera is governed by a kind of platform logic that demands that seeing and sensing are carried out in a distributed fashion, requiring the equal formatting of each component connected within.

As Mackenzie and Munster affirm, smartphones do not create representational images but are ‘entire sensing “platform[s]” capable of carrying out the distribution and integration of different forms of processing’ (Mackenzie & Munster, 2019, p. 14). As they further contend, a smartphone’s image sensor ‘is itself already a mini-workstation for image processing’ (Mackenzie & Munster, 2019, p. 14). In ‘recent high resolution digital cameras, sensors are panchromatic and arranged in an array’ (Mackenzie & Munster, 2019, p. 14). In such a case, the sensor ‘does not merely *receive* light but *process* light quantities alongside or in tandem with other information’ (Mackenzie & Munster, 2019, p. 14, emphasis added). As Mackenzie and Munster go on to write, smartphone image signal processors (ISPs) that manage smartphone camera-related problems (such as motion blur and low-light) are merely a ‘downsized iteration of ... image recognition processors for autonomous vehicles’ (Mackenzie & Munster, 2019, p. 16).

Beyond the distributed nature of sensing, Mackenzie and Munster also consider developments in computation itself. Only through the development of graphics processing units (GPUs), historically used for

processing computer game graphics, are smartphones able to ‘render images aggregately computable through massive *calculative parallelism*’ (Mackenzie & Munster, 2019, p. 17, emphasis added). This calculative parallelism enables ‘vast numbers of discrete arithmetic operations’ to be ‘carried out in parallel lanes’ (Mackenzie & Munster, 2019, p. 17) in order to generate images. Not limited to the world of smartphone cameras, finding computationally efficient ways to process huge volumes of sensor data is central to the work of machine vision researchers in the domain of autonomous driving.

Perhaps missing from Mackenzie and Munster’s account is an articulation of the ‘contact points’ between operative elements. That what they call platform seeing is enabled through certain kinds, and levels, of interoperability. In addition, that the ‘distribution’ and ‘redistribution’ of sensing have a certain specificity, rather than simply involving a new mix, or rearrangement, of agents involved in this sensing. *Where*, in other words, is sensing distributed to? To *what* is sensing redistributed to? *When* is sensing distributed or redistributed? In the development of autonomous vehicles, these questions rarely receive the same answer. Sensing may rely on more ‘sovereign’ (Hind, 2022) systems such as lidar or engage other such sensing systems (such as radar) at certain distances. Moreover, that this interoperability—or, distributed sensing capacity—must be shaped, managed, optimized, and ‘finessed’ in order to be made (at least provisionally) operational. It must, in other words, be made to work.

As I have also suggested elsewhere, ‘interoperationality precedes operability, rather than vice versa’ (Hind, 2023, p. 4). This means that the specific operation of certain systems does not occur before any subsequent interoperation with connected systems. Instead, ‘the interoperability and interoperationality of each unit of sensor data must be ensured before any one system “operates”’ (Hind, 2023, p. 4). As a result, interoperability—as a matter of concern and as a practical task—must precede operability. Thus, ‘operability is dependent firstly on interoperability’ (Hind, 2023, p. 4).

Practising Interoperability

Accordingly, interoperability manifests differently in different cases, and thus, interoperation is not simply a binary proposition. Instead, interoperation is performed in various ways, with varying degrees of success. Yet, in all cases, systems dependent on interoperation are always more than the sum of their parts.

Thus, there is a *technical interoperability*, in which technical systems such as systems for generating point cloud data (lidar) or video footage (cameras) must be made to work together, or with other software, to make sense of raw data. This technical interoperability is what Distelmayer (2018) refers to as interrelations or thresholds: between hardware and software, computers and more computers, or computers and non-computer forms. Yet, I argue that interoperations are different from interrelations or interfaces, in that these entities (and the data that flows through, and enlivens, them) are dependent on each other to function, and therefore deficient, or partial, without interoperation.

But there is also, following Wilmott (2016, 2020), an *epistemic interoperability*, requiring ‘discursive and linguistic compatibility in order to work’ (Wilmott, 2016, p. 8) or what Distelmayer (2018, p. 62, authors’ emphasis) again would call the interrelations between humans and computers, concerning ‘[o]perations as *us* dealing with *them*’. For Wilmott, technical interoperability alone (such as a key card to open a gate) is not enough. Instead, technical interoperability also requires this epistemic interoperability which at some level (i.e. a discursive one) must be made interoperable, just like on a technical level. For instance, in how ‘the relationships between Spatial Big Data and cartographic reason as interoperable discursivities and logics enabled an ever-expanded ordering of spatial knowledge’ (Wilmott, 2016, p. 2), constituting a ‘desire to create interoperable systems (compatible systems) and increasingly universal narratives in universal languages’ (Wilmott, 2020, p. 20).

This claim to universality—in ‘360° sensing’, for instance—is routinely, and often casually, referred to by autonomous vehicle manufacturers (Oxa, 2019), despite the modular, interoperable nature of both the technology and attendant discourses. As the specific processes hopefully

show, later, interoperability is dependent on an ‘unsettled relation’ (Azar et al., 2020, p. 5) between a technical dimension (‘the techne to make an image’ [Azar et al., 2020, p. 5]) and this discursive dimension (‘the socio-cultural milieu that allows for certain technics, and images, to emerge’ [Azar et al., 2020, p. 5]), as various traditions within machine vision and computer science are mobilized. It is what Adelheid Voskuhl (2004, p. 415) has referred to in another context as ‘functional contingency’, in which the work done to achieve interoperability assumes a somewhat precarious, but nonetheless operational, form.

Thus, in the context of autonomous vehicles specifically, something like a situational interoperability exists, in which technical and epistemic interoperability is entwined. Here, the interoperability is not referencing a fundamental, technical interoperability or an off-the-shelf, plug-and-play compatibility, but a continual, precarious practice (interoperation) that requires that two or more systems to interoperate. This shift from interoperability as (intended, desired, or actual) state to interoperation as continual practice allows one to study interoperation *in action* as it happens.

As Hoel (2018) intimates, there are many working theories about ‘operations’ within technical disciplines like computer science. This is no less the case within the world of machine vision and, even more specifically, within the domain of autonomous vehicles. Here, ‘the history of the development of autonomous vehicle, stretching back to Stanley [DARPA Grand Challenge 2005 winner], Uber ATG, and Cruise, suggests that autonomous driving demands *thinking* and *acting* in an interoperational fashion’ (Hind, 2023, p. 4, authors’ emphasis).

Yet, whilst there is an acknowledgement that such work demands interoperability, leading autonomous vehicle engineers who spoke at the CVPR workshop (WAD) contended that much of this work takes place in a far more modular, silo-d, and sequential fashion. In her keynote speech, Raquel Urtasun, previous member of the KITTI team, and former Uber ATG chief scientist, spoke about her new autonomous vehicle project, Waabi (Urtasun, 2021). Carl Wellington, previously Perception Lead at Uber ATG and at Carnegie Mellon University’s (CMU) famous Robotics Institute, spoke about his work as Head of Autonomy at self-driving start-up, Aurora (Wellington, 2021).




Fig. 4.1 A simple diagram of perception, forecasting, and planning. (Source: Wellington, 2021)

In different ways, both Urtasun and Wellington acknowledged this blind spot in contemporary software development. For Urtasun, the extant problem is that a ‘traditional’ approach to self-driving is executed through the ‘autonomy stack’ in which different modules are responsible for different tasks (Fig. 4.1)—modules that this book analytically replicates to some degree, of mapping, perception, prediction, planning, and control. As Urtasun (2021) suggests, these modules ‘are not trained for end-task’ but for their respective module-dependent tasks. Accordingly, it means that this autonomy stack that Urtasun (2021) refers to is ‘developed in silos’.

The result is that the ‘small interfaces between modules results in cascading errors’ (Urtasun, 2021) as one moves through this software stack—evidently a problem of interoperability between these modules or levels in the stack itself. Thus, ‘any technical solutions or methods devised to resolve problems related to the operation of these discrete [modules]’ (Hind, 2023, p. 4) necessarily involve building ‘more and more modules’ (Urtasun, 2021). In this, any ‘holistic view’ to delivering autonomous driving is deemed impossible, thanks to the entrenchment and familiarity of the module/stack approach.

Urtasun’s solution is an ‘end-to-end’ approach in which a ‘single AI system’ is developed, with ‘all modules trained for the end task’ (Urtasun, 2021) rather than each module (mapping, perception, prediction, etc.) being training for specific modular tasks. Skirting around many of the connected problems with developing such an approach, Urtasun clearly suggests that interoperability is a necessary requirement for ‘solving’ autonomous driving.

Wellington thinks similarly, considering how his Aurora team might combine different modules, in the hope of smoothing (or wholesale removing) the interface between them. In particular, Wellington

introduces ways to combine perception and motion planning, with so-called actor forecasting (Wellington, 2021) of the future states of road users in-between. Part of a reorganization of this work, as Wellington lays out, might involve ‘thinking about the perception forecasting problem as a single problem’ (Wellington, 2021), rather than as two distinct, nonetheless interconnected, problems. This is what Wellington straightforwardly refers to as ‘joint perception and forecasting’ (Fig. 4.2), in which a combined model is built incorporating both stages of the traditional autonomy stack.

What is interesting is that Wellington, in addition to what he calls the ‘left to right perspective’—joint, interoperable work done from the more foundational level (perception) to a higher level (forecasting)—one might also work from ‘right to left’ (Wellington, 2021). That is, by combining forecasting and planning modules instead. This ‘conditional forecasting’, as Wellington (2021) describes, results in decisions being made by the autonomous vehicle being fed or ‘passed back’ through the model to inform the forecasting work being undertaken. In both cases, Wellington is proposing innovative solutions to a lack of interoperability within the ‘autonomy stack’ as Urtasun refers to it.

The source of these many of the interoperational problems identified by both Urtasun and Wellington lies in the richness of sensor data being collected on the move by, and in, a vehicle. The necessary, and variable, movement of an autonomous vehicle leads sensors (lidar, cameras, etc.) to endlessly capture new scenes requiring processing and begetting inter-operation. At one moment, for instance, a vehicle may be moving along at 30 km/h, before accelerating to 50 km/h, and then slowing to 20 km/h. Throughout, the systems on-board an autonomous vehicle must be equipped to cope with this constant stream of sensor data being




Fig. 4.2 Joint perception and forecasting. (Source: Wellington, 2021)

generated in, and through, movement, but also to process it quickly, efficiently, and accurately in order to make critical control decisions such as to brake or swerve. Thus, mobility generates the sensor data necessary for further movement. However, it is only through interoperation that perceptive issues arising from the capture of the data can be resolved, and without which mobility cannot be maintained.

Finessing Interoperation

To avoid making a sharp distinction between different kinds, or levels, of interoperability, I draw on the work of anthropologist Michael Fisch. In his ethnography of the Tokyo commuter train network, Fisch (2018, p. 40) talks about ‘finessed interoperability’ in which *excess* and *leeway* are generated in such a system. He discusses how commuter train operators engage in so-called recovery driving (Fisch, 2018, p. 40) to make up for lost time mandated by allotted station ‘dwell time’ trains are expected to wait at each stop for. Here, interoperability is not considered as a strictly technical relation in which two devices or technologies speak to each other alone, but the way in which one technology (the commuter train) is operated in accordance with the rules and demands of the another (a schematic traffic plan-diagram known as a *daiya* in Japan). ‘For train drivers’, as Fisch (2018, p. 40) contends, ‘producing *joyū* [leeway, space] demands a unique combination of intuition, technique, and attunement to the shifting conditions of operation’. Interoperability, then, is a kind of learned or programmed relation in which each interoperable component (machinic, human, or otherwise) becomes attuned and *in sync* with another. This interoperability is only made possible through a contingent settling of the relation between technical and discursive elements.

The terms interoperability and interoperation are a way to straddle, and make sense of, typically distinct concerns around sensing, decision-making, and algorithmic systems. In doing so, the intention is similarly to extract the meaning of relevant terminology of computer scientists and machine vision experts, to afford a higher-level view of these motivations for greater interoperability, as voiced by the likes of Urtasun and Wellington. Whilst they each, at first, appear as technical issues to be

settled or resolved (or at least optimized), in being re-articulated as examples of interoperation-in-action, they instead appear as ‘sensor strategies’ in which the interoperability of relevant systems is *finessed* by interest parties, such as machine vision experts. That, in other words, it is through the finessed optimization of these processes (3D object detection, streaming perception), autonomous vehicles emerge as particular socio-technical objects capable of sensing, and making sense of, the world.

These sensor strategies cannot, I argue, necessarily be seen as forms of AI ‘micro-work’ increasingly common in automotive settings (Hind, 2021). None can be considered as comprising ‘small, fragmented tasks performed remotely online’, nor involve ‘mundane, repetitive, [or] atomized’ labour (Tubaro & Casilli, 2019, p. 334) common to most forms of AI micro-work. To some degree, however, the work discussed here can be considered both reliant on such micro-work, say, in the labelling of images discussed in Chap. 3 but also ‘downstream’ of such work, in that they are nonetheless in some operational relationship, merely embedded at different stages, and in different locations, in the autonomous vehicle development process. As will be discussed later, the practice of ‘finessing’ interoperation hinges on the relative (professional) possibilities that machine vision researchers possess to explore solutions to extant computational problems. This practical latitude for problem-solving is not one common to most forms of micro-work which, instead, generally consist of strictly structured, and monitored, work practices. For instance, in how micro-workers are expected to categorize objects in a photograph according to formal, predetermined criteria. By contrast, as Fisch (2018) discusses, the art of finessing something (an act, a task, etc.) involves both skill and dexterity, as well as creative instinct and flair: what Voskuhl (2004, p. 406) has called ‘tinkering’ strategies, and what Agre (1997) understands as a tenet of AI work in general. As machine vision researchers working in AI research centres or automotive research and development (R&D) facilities, such opportunities are typically more open to them.

This work is what Rieder refers to as ‘algorithmic techniques’ (Rieder, 2020), where different ‘habitats’ offer the opportunity for some techniques to ‘thrive’ rather than ‘whither’ (Rieder, 2020, p. 247). Here, the general goal is to ‘design a system that produces “good” results in the

domain of its application' (Rieder, 2020, p. 252), with ML and machine vision processes explicitly developed to 'produce operationally viable results rather than scientific models' (Rieder, 2020, p. 257) whether for language, as Rieder writes, or for the general interpretation of images and objects. As Rieder understands it, such work functions as a 'trading zone' following Galison (1996), where 'statistics and other areas of mathematics intermingle with ideas about language, information and knowledge as well as computing machinery, systems design, and the concrete and imaginary requirements of "knowledge workers" and "decision-makers"' (Rieder, 2020, p. 258). Here, the habitat for such work is necessarily a hybrid one, where the finessing of inoperability is enabled through the setting of the conditions of the habitat in the first instance—of what can, and cannot, be done. As Rieder concludes, 'knowing *why* a technique works well is not a fundamental requirement' (Rieder, 2020, p. 259, authors' emphasis), only that it does, in fact, work.

As Adam Sargent et al. (2021, p. 565) discuss, the finding of solutions to operational problems within such a context 'highlights the complex human-machine configurations through which defects are perceived, identified and interpreted'. Professional engineers, in Sargent et al.'s case, in the US steel industry are engaged in 'perceptual work' involving 'sensing defects', that does not simply correspond to the 'internal processes of the engineer' (Sargent et al., 2021, p. 565) but to the assemblage of sensing technologies at their disposal that allows them to sense defects in steelwork and the steel production process. In the cases to follow, I emphasize how similar such perceptual work is performed, focused on the identification of operational errors (rather than defects, *per se*) in the machine vision-dependent processes. Here, errors or faults do not present themselves as things to be 'fixed' or successfully, and completely, eradicated. Instead, such calculative errors can only ever be resolved in relation to any 'desired outcome' functioning as 'the central locus of normativity' as Rieder (2020, p. 252) writes.

Moreover, that there are perhaps two levels, or orders, of perceptual work being undertaken here. Firstly, that the work concerns machine vision and algorithmic systems capable of machine vision. In this, such systems obviously engage in their own 'perceptual work', recognizing and categorizing objects, segmenting lidar points, or generating path

trajectories of other road users. However, in much the same way in which the sensing of defects in the steel production process does not simply correspond to the ‘internal processes of the engineer’ (Sargent et al., 2021, p. 565), neither does the sensing of potential risks on the road simply correspond to the ‘internal processes’ of the system. Instead, various parameters are played with, set, and programmed by the machine vision researchers, as Chap. 5 contends. Secondly, the machine vision work being undertaken also requires interpretive, evaluative, analytical work to make sense of the decisions being made by the perception system itself. It is this level that corresponds to the perceptual work discussed by Sargent et al. (2021), nonetheless considerably shaped by the perceptual work undertaken by the perception system, and the relative agency and autonomy it has in performing its own work. It is arguably another example of the ‘unsettled relation’ (Azar et al., 2020, p. 5) between the technical and the discursive.

I now discuss two sensing processes that have required the invention of particular ‘sensor strategies’ to facilitate the intended interoperability of autonomous vehicles: 3D object detection and streaming perception. As outlined in the introduction, I encountered both of these techniques during participation in two machine vision events: a virtual summit on computer vision held in June 2020 called ‘Machines Can See’, and a ‘Workshop on Autonomous Driving’ (WAD) held in June 2021, as part of the Computer Vision and Pattern Recognition (CVPR) conference. It is in these kinds of events where ‘cutting-edge methods, techniques, and approaches are shared with those working in related fields of computer vision and deep learning’ (Hind, 2023, p. 5), constituting important settings where sensor strategies are shared with members of the wider machine vision community.

3D Object Detection (Visibility Volumes)

3D object detection involves the capacity to detect and orient vehicles, and other road users, within 3D space. It is a critical capacity of all autonomous vehicle systems. Lidar is the most commonly used sensor technology to enable this detection work to take place. Despite its many

advantages, lidar has one significant disadvantage: lidar points effectively destroy or omit data on phenomena behind each lidar point captured. Upon hitting an object, a lidar point is returned, meaning any secondary object hidden behind this initial object is not captured at all. The result is that the object from the capture scene is not represented in the data—essentially destroyed within the dataset being generated.

As Hu et al. (2020, p. 2) write, ‘once a particular scene element is measured at a particular depth, visibility ensures that all other scene elements behind its line-of-sight are occluded’. As a result of this loss of data, ‘such 3D sensed data might be better characterized as “2.5D”’ (Hu et al., 2020, p. 2). Naturally, hampered by lidar’s ‘occlusion’ capacities, any such autonomous vehicle making use of lidar ends up capturing any given scene in much lower fidelity (Fig. 4.3).

Accordingly, lidar is rendered as a far less interoperable medium as might otherwise be thought. Rather than being understood as the best of all available sensing options, its weaknesses as a data capture method are exposed. In such a case, lidar is shown to be an imperfect representation, not only failing to capture occluded objects but expunging them from the record altogether. This has significant repercussions for interoperability. Any object-recognition system designed to ingest lidar data without any such remedy for its ‘2.5D’ nature risks the total ignorance of real-world objects not translated into their lidar-dependent worlds. Whilst such objects might end up in the latter in subsequent frames (say as a car moves into shot), any such system would still be at an operational disadvantage, only capturing the existence of occluded objects at their moment




Fig. 4.3 Examples of different 3D sensor data representations with ‘freespace’ visualized (right). (Source: Hu et al., 2020)

of non-occlusion. Without such caution being taken, the vehicle runs the risk of reacting too slowly (or not at all) to objects moving from their occluded positions to visible ones.

Hu et al. (2020)—researchers at Argo AI in Pittsburgh, US—offer a solution referred to as ‘raycasting’. When a lidar point hits an object, it records a coordinate. When repeated in a so-called lidar ‘sweep’, a series of coordinates are recorded. In between the many lidar points, there is ‘freespace’, where no such data is recorded at all. In generating a ‘3D voxel grid’ in which each coordinate is recorded as either ‘occupied, free, or unknown’ (Hu et al., 2020, p. 4), each coordinate within 3D space can be provided with a known value, regardless of whether a lidar point has hit an object or not. Referring to the output as a ‘visibility volume’ (Hu et al., 2020, p. 4), Hu et al.’s method can be considered a technique for producing fully 3D data—a remedy for the ‘2.5D’ lidar data originally captured.

Rendering lidar data in ‘full’ 3D form, assisted by the 3D voxel grid, offers a stronger form of interoperability between the original lidar input and a desired feature map output, required in the ML model process. Lacking a 3D voxel grid, lidar generates a coarse view of the world, able to capture the presence of some objects, but unable to make sense of occluded objects. In adding interstitial stages, Hu et al. are able to enhance and *augment* the perceptive qualities of lidar. In so doing, they produce a ‘diagrammatic abstraction’ (Mackenzie, 2017, p. 55) eminently more useful for the ML-driven system they are building.

Returning to the question of how this affects the operational pipeline, having a record of what the researchers call a ‘visibility volume’ (in addition to lidar data) enhances the 3D object-recognition process. Using their approach on the NuScenes 3D detection dataset, a commonly used autonomous vehicle dataset, they achieve a mean average precision (mAP) score improvement of 4.5%—a standard metric for evaluating the performance of ML-driven object-recognition processes (Hu et al., 2020). In particular, their model offers big improvements on detecting heavily occluded cars (0–40% visible), with significant implications for how an autonomous vehicle would react to a given scenario involving such road users.

Understanding this conceptually, *augmenting* operability offers a relatively straightforward way of ensuring wider interoperability. Indeed, this augmentation is itself already a fusion of kind at least in the second stage of the approach, as lidar point sweeps are combined by Hu et al. (2020) with the visibility volumes generated through their creation of a 3D voxel grid. Whilst these additions undoubtedly add greater complexity into the object detection/recognition process, they do so in order to properly account for deficiencies of lidar that will only be encountered later in the pipeline during forecasting and planning. Resolving these latent issues at this stage in the process only serves to strengthen later interoperability down the line.

Streaming Perception (Dynamic Scheduling)

Streaming perception¹ concerns the processing of video frames and the ‘algorithmic trade-off’ (Li et al., 2020, p. 1) between *accurate* image understanding and *quick* image understanding. Accuracy in such cases can be defined by achieving a certain threshold of objects correctly identified and categorized, and speed defined by the completion of a perception process in advance of subsequent phases of operation, such as motion forecasting and planning. A streaming perception process that is inaccurate or too slow risks generating errors along the operational pipeline of an autonomous vehicle model. Managing this algorithmic trade-off between accuracy and speed is thus of great importance, despite the extant problems associated with doing so appropriately. The main computational limitation in such a use case as autonomous driving concerns the ability for any such system to process a volume of sensor data, accurately, at speed, or vice versa.

For example, if speed is privileged over accuracy in such a process, a cyclist might end up being mis-categorized as an ordinary vehicle and deemed by an autonomous vehicle to be moving at the (faster) speed of a vehicle rather than a bicycle. As a result, the cyclist might be at a higher risk of being hit by the autonomous vehicle, with the latter believing the

¹ Various referred to as streaming processing optimization and streaming image understanding by Li et al. (2020).

cyclist was travelling faster, and therefore out of the path of the oncoming autonomous vehicle. If accuracy is privileged over speed, whilst the cyclist in such a scenario might be correctly categorized, the result of the decision might be communicated too late to affect the execution of the manoeuvre needed to take account of them.

One of the solutions devised again by researchers at Argo AI in order to optimize streaming perception is something called ‘dynamic scheduling’ (Li et al., 2020, p. 1). Dynamic scheduling skips so-called stale frames that might otherwise be fully processed and evaluated under usual circumstances, even if these frames no longer correspond to the current, real-world state of the vehicle. Under these usual circumstances, the system continues busying itself with analysing video frames *regardless* of whether they correspond to the real-world state of the vehicle, whether stale or ‘fresh’. Outside of this narrow computational concern, staleness matters. Processing video frames that no longer correspond to the state of the environment is evidently pointless, with each exact snapshot capturing a moment lost in time. They are, so to speak, already in the rear-view mirror, images of other road users in past states, performing past manoeuvres.

In instructing the system to skip the evaluation of stale frames, dynamic scheduling readies the system for tackling imminent fresh frames, those that *do* correspond to the real-world state of the vehicle and therefore to any emergent risks that might present themselves on the road. In this, dynamic scheduling enables a certain level of breathing space, or computational capacity, unburdened with a never-ending stack of frames it is otherwise instructed to evaluate. As a result, the system is better prepared for fresh frames that *do* matter: images of other road users in *current* states, performing *current* manoeuvres, as well as possibly *imminent* ones too.

Yet, as Li et al. (2020, p. 2, emphasis added) write, ‘*latency* is inevitable in a real-world perception system’. In the above description of an ordinary scheduling procedure, it is not that the system is purposely or knowingly evaluating stale frames, just that when it ‘takes a snapshot of the world at t_1 ... when the algorithm finishes processing this observation, the surrounding world has already changed at t_2 ’ (Fig. 4.4). The question for Li and his colleagues, then, is how to *optimize* the processing of




Fig. 4.4 A simplification of the streaming perception process. (Source: Li et al., 2020)

streaming images such that the system is best able to comprehend any *current* situation (point B in Fig. 4.4), rather than any *previous* situation (point A in Fig. 4.4). In practical terms, this latency problem renders the vehicle ill-equipped to act but also *react*. As Li et al. (2020, p. 1) contend, ‘[a] crucial quantity governing the responsiveness of [an autonomous] agent is its reaction time’.

What Li et al.’s solution or sensor strategy involves, then, is the *parallel processing* of frames, such that the system is able to tackle any new (‘fresh’) frame at the same time as any existing-going-stale frame. This is a version of what MacKenzie and Munster (2019, p. 17) have referred to as ‘calculative parallelism’. Any ordinary scheduling arrangement would simply queue each frame up sequentially until the system is ‘finished’ evaluating one, before moving on to the next. In this, dynamic scheduling involves a kind of computational multi-tasking, in which the system is designed to juggle the evaluation of more than one frame at a time. That is, rather than prioritizing sequential completion. The main operating principle behind dynamic scheduling is thus that ‘streaming perception requires understanding the state of the world at all time instants’ (Li et al., 2020, p. 2), such that the overriding aim of a system engaged in such perception should be ‘to produce accurate state estimations in a timely manner’ (Li

et al., 2020, p. 11). In other words, that if and when faced with other tasks (i.e. the processing of point A as in Fig. 4.4), the system must always prioritize the *current situation*; even, or especially, at the expense of ‘completing’ the processing of a prior frame in a sequence. Here, ‘completion’ is not considered a priority, less a distraction from the evaluation of the current state of the driving environment. Instead, using what Li et al. (2020), p. 11) refer to as a ‘shrinking-tail policy’, an algorithm following a dynamic schedule is happy to ‘sit idle and wait’ (Li et al., 2020, p. 11) if the next available frame is already stale.

Dynamic scheduling is intended to smooth the interoperation between video input and classification outputs, meaning that critical processing capacity is not wasted by parsing ‘useless’ frames (Fig. 4.5). Underpinning this work is an implicit understanding that any delays or errors introduced during the perception phase will have knock-on effects, undoubtedly resurfacing at either forecasting or planning stages if not at the point of vehicle control itself. What dynamic scheduling offers, therefore, is a re-constitution or calculation of the necessary *ordering* and *pacing* of interoperability. Rather than requiring some form of augmentation to facilitate interoperability between different phases of operation, like through the generation of visibility volumes, here the question is of when the *same* tasks might be best performed to facilitate interoperability, not whether *additional* tasks should be. Here, interoperability does not therefore necessarily just depend on adding or attaching things onto existing processes (or indeed, removing them entirely) but by re-designing their processual execution, relevant to the current situation.



Fig. 4.5 Dynamic scheduling. (Source: Li et al., 2020)

Finessing Scheduling

In this final part, I want to draw a further comparison between Li et al.'s articulation of the scheduling of video frames and Fisch's (Fisch, 2018) discussion of the scheduling of trains on the Tokyo commuter train network. Both concern methods to deal with systems that involve the management of things (frames, trains) that must be 'processed' quickly and efficiently. For system operators of Tokyo's commuter train network, this involves managing the gap between a 'painstakingly calculated, idealized' schedule designed by professional rail technicians known as a 'principal *daiya*' and an 'actual, "operational" (*jisshi*) *daiya*' corresponding to the 'lived tempo of the city and train network' (Fisch, 2018, p. 5). Together, as Fisch explains, the *daiya* assumes a 'dynamic quality' (Fisch, 2018, p. 5) as system operators seek to address the gap between the idealized and actual schedule. On Tokyo's commuter train network, however, the task is not to reduce this gap—itself an impossible job as 'operation beyond capacity' is infamously the norm, as Fisch (2018, p. 5) explains—but to actively manage it. For drivers, this involves '*fnessing the interval* within the commuter train network's margin of indeterminacy' (Fisch, 2018, p. 29, authors' emphasis).

What seems evident within Li et al.'s work on dynamic scheduling is that something similar is at play, albeit with the explicit involvement of a technical actor. In this case, the machine vision experts play the role of the rail technicians, setting the parameters for an idealized scheduling of video frames. In place of Fisch's commuter train drivers, however, is an algorithmic system, instructed to perform a set of evaluative tasks in a non-sequential, 'dynamic' manner. Here, the margin of indeterminacy of the Tokyo commuter train network with which drivers must learn to deal with is replaced in Li et al.'s work with a kind of *evaluative incompleteness*. The algorithmic system must likewise learn to de-prioritize the processing of frames *completely*, and instead value *liveness*, the lived tempo of the autonomous vehicle in its environment (Fig. 4.6). Li et al.'s system is expected to operate 'on time' in ways that Tokyo's commuter train drivers are. Just as the latter deploy a number of "'speed up" strategies' (Fisch, 2018, p. 35), so Li et al. offer suggestions as to how the typical stages of

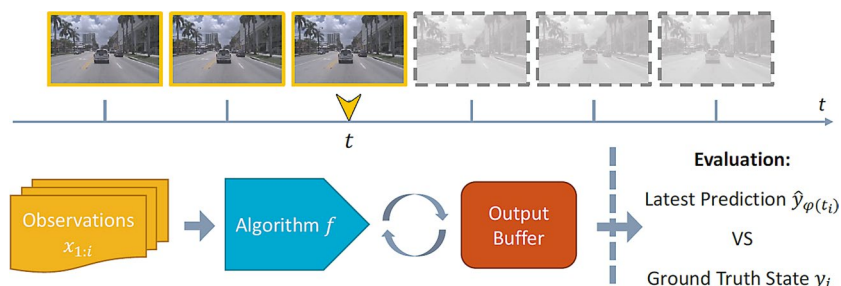


Fig. 4.6 Processing video footage from an autonomous vehicle. (Source: Li et al., 2020)

streaming perception (detection, association, forecasting) might be better *integrated*, *ordered*, or *blurred* to speed up the processing of frames.

Here, I argue that in emphasizing *integration* and the *blurring* of module *boundaries*, Li et al. (2020) are principally concerned with, as well as trying to resolve, the question of interoperability. This interoperability is being tackled both at the level of modules *within* algorithmic processes (detection, tracking, and forecasting stages), but also *between* (concurrent, infinite) GPUs, and necessarily between algorithm and the sensor system feeding sensor data (in the form of static video frames) into the algorithmic process. The various techniques employed by Li et al. (multiple GPUs, module blurring, non-sequentiality) that together comprise what they call dynamic scheduling can be seen as a specific sensor strategy that deals—or attempts to deal with—the distributed perceptive capacities of autonomous vehicles, what MacKenzie and Munster (2019) call platform seeing.

As the above account hopefully shows, these various techniques involve the finessing of interoperability. As Fisch (2018, p. 31) explains:

To finesse something is to make it work when, logically, speaking, it should not. Finesse is about pulling something off against all odds. Invoking terms like *flair*, *panache*, or *élan*, finesse bespeaks a method irreducible to skill, expertise, or systematicity. Finesse transcends the logic of rational methods whereby cause and effect can be situated as calculable corollaries; it involves instead qualities like instinct, affect, and feeling – qualities that are embodied, sensual, and informed by the precarious order of contextual relations ...

More importantly, it suggests a relation that transpires as a kind of dialogue in the mode of technicity between provisionally stable processes rather than established and fixed ontologies.

In this, Li et al. are feeling their way towards a plausible, practical solution ‘good enough’ for the task-at-hand, or indeed, ‘better than’ an established alternative. In this there are no logical ‘fixes’ that settle extant problems such as ‘how autonomous vehicles see’, only a more modest movement towards streaming perception itself.

Conclusion

The cases I have explored in this chapter hopefully evidence what interoperation looks like in the wild, in respect to machine vision, in which interoperation is deemed necessary in order to resolve problems that arise with sensing itself. These are problems, I argue, that are not answered by *not* sensing or necessarily reducing reliance on sensing but are only be fixed through building complimentary systems such that they are made interoperable with these extra systems that work to identify, correct, and omit gaps and errors. All these cases have centred on the limitations of different sensing formats (video, RGB image, lidar), each themselves unable to process the data they capture and thus unable to address the limitations of their format alone, thus necessitating interoperation with systems that can and do. I have referred to the techniques that engineers use to resolve these questions as ‘sensor strategies’ in which they aim to *finesse* the interoperation of machinic systems.

In this, machine vision researchers working in the domain of autonomous driving are enrolled within, and speak to, the wider machine vision community, as *problem-solvers* in which their efforts are (narrowly) oriented towards the finding of technical solutions to ‘tricky problems’ that otherwise are meant to *improve* the sensing capacities of possible autonomous vehicles. In this, such problems are *practically resolvable* or *operationally resolvable* in that solutions can be ‘more-or-less’ found that could be, or are, acceptable to autonomous vehicle firms.

Interoperability, I have argued, is integral to the comparative success of such efforts in which additional optimization, processing, detection, segmentation, and augmentation are performed on sensor data generated by, and in, lidar point clouds and video streams. In the cases I have explored, this ‘substrate’ sensor data is considered valuable, that is ‘operational’, only through secondary optimization, detection, or processing, without which the sensor data falls below acceptable thresholds for usability.

Yet, as I have suggested, the cases I have described do not result in ‘perfect’ solutions in which the sensing capacities are now deemed ‘accurate’ or ‘perfect’. Indeed, all the cases discussed involved significant inter-operational ‘trade-offs’ that come with integrating such secondary processes. These trade-offs are part-and-parcel of designing sensor systems for autonomous vehicles, in which sensor data is generated ‘on the move’ imposing computational constraints on how well or completely sensing can be performed.

References

- Agre, P. (1997). Toward a critical technical practice: Lessons learned in trying to reform AI. In G. C. Bowker, S. Leigh Star, W. Turner, & L. Gasser (Eds.), *Bridging the great divide: Social science, technical systems, and cooperative work* (pp. 131–158). Lawrence Erlbaum Associates.
- Amoore, L. (2018). Cloud geographies: Computing, data, sovereignty. *Progress in Human Geography*, 42(1), 4–24. <https://doi.org/10.1177/0309132516662147>
- Apple. (2023). iPhone 14 Pro – technical specifications. *Apple*. Retrieved October 1, 2023 <https://support.apple.com/kb/SP875>
- Azar, M., Cox, G., & Impett, L. (2020). Introduction: Ways of machine seeing. *AI & Society*, 36, 1093–1104. <https://doi.org/10.1007/s00146-020-01124-6>
- Bunz, M. (2019). The calculation of meaning: On the misunderstanding of new artificial intelligence as culture. *Culture, Theory and Critique*, 60(3–4), 264–278. <https://doi.org/10.1080/14735784.2019.1667255>
- Distelmayer, J. (2018). Carrying computerization: Interfaces, operations, deprecations. In L. Feiersinger, K. Friedrich, & M. Queisner (Eds.), *Image – action – space: Situating the Screen in visual practice* (pp. 55–68). De Gruyter. <https://doi.org/10.1515/9783110464979-005>
- Farocki, H. (2004). Phantom images. *Public*, 29, 12–22.

- Fisch, M. (2018). *An anthropology of the machine: Tokyo's commuter train network*. The University of Chicago Press.
- Galison, P. (1996). Computer simulations and the trading zone. In P. Galison & D. J. Stump (Eds.), *The Disunity of science: Boundaries, contexts, and power* (pp. 118–157). Stanford University Press.
- Halpern, O. (2014). *Beautiful data: A history of vision and reason since 1945*. Duke University Press.
- Hind, S. (2019). Digital navigation and the driving-machine: Supervision, calculation, optimization, and recognition. *Mobilities*, 14(4), 401–417. <https://doi.org/10.1080/17450101.2019.1569581>
- Hind, S. (2021). Making decisions: The normal interventions of Nissan 'mobility managers'. *Mobilities*, 17(4), 467–484. <https://doi.org/10.1080/17450101.2021.1988682>
- Hind, S. (2022). Machinic sensemaking in the streets: More-than-lidar in autonomous vehicles. In G. Rose (Ed.), *Seeing the city digitally: Processing urban space and time* (pp. 57–80). Amsterdam University Press. <https://doi.org/10.2307/j.ctv2j6xrs3.6>
- Hind, S. (2023). Sensor work: Enabling the interoperation of autonomous vehicles. *Journal of Computer-Mediated Communication*, 28(5), 1–10. <https://doi.org/10.1093/jcmc/zmad014>
- Hoel, A. S. (2018). Operative images: Inroads to a new paradigm of media theory. In L. Feiersinger, K. Friedrich, & M. Queisner (Eds.), *Image – action – space: Situating the screen in visual practice* (pp. 11–22). De Gruyter. <https://doi.org/10.1515/9783110464979-002>
- Hu, P., Ziglar, J., Held, D., & Ramanan, D. (2020). What you see is what you get: Exploiting visibility for 3D object detection. Preprint. <https://arxiv.org/abs/1912.04986>
- Li, M., Wang, Y.-X., & Ramanan, D. (2020). Towards streaming perception. Preprint. <https://arxiv.org/abs/2005.10420>
- Mackenzie, A. (2017). *Machine learners: Archaeology of a data practice*. MIT Press.
- Mackenzie, A., & Munster, A. (2019). Platform seeing: Image ensembles and their invisualities. *Theory, Culture & Society*, 36(5), 3–22. <https://doi.org/10.1177/0263276419847508>
- Oxa. (2019, November 8). Our eyes are a vital asset when it comes to spotting and reacting to danger on the roads. *Twitter*. https://twitter.com/Oxa_UA/status/1192815056926457858
- Parisi, L. (2020). Negative optics in vision machines. *AI & Society*, 36, 1281–1293. <https://doi.org/10.1007/s00146-020-01096-7>

- Rieder, B. (2020). *Engines of order: A mechanology of algorithmic techniques*. Amsterdam University Press. <https://doi.org/10.1515/9789048537419>
- Sargent, A., Vinson, A. H., & Stevens, R. (2021). Sensing defects: Collaborative seeing in engineering work. *Social Studies of Science*, 51(4), 564–582. <https://doi.org/10.1177/0306312721991919>
- Tubaro, P., & Casilli, A. A. (2019). Micro-work, artificial intelligence and the automotive industry. *Journal of Industrial and Business Economics*, 46, 333–345. <https://doi.org/10.1007/s40812-019-00121-1>
- Uliasz, R. (2020). Seeing like an algorithm: Operative images and emergent subjects. *AI & Society*, 36, 1–9. <https://doi.org/10.1007/s00146-020-01067-y>
- Urtasun, R. (2021, June 20). CVPR '21 WAD keynote – Raquel Urtasun, Waabi/University of Toronto. *YouTube*. <https://www.youtube.com/watch?v=PSZ2Px9PrHg>
- Velodyne. (2017, January 2). It began with a race...16 years of Velodyne Lidar. *Velodyne*. <https://velodynelidar.com/blog/it-began-with-a-race16-years-of-velodyne-lidar/>
- Velodyne. (2019). *User's manual and programming guide: HDL-64E S3*. Velodyne.
- Velodyne. (2021, May 19). Velodyne's HDL-64E lidar sensor looks back on a legendary career. *Velodyne*. <https://velodynelidar.com/blog/hdl-64e-lidar-sensor-retires/>
- Volmar, A., Jancovic, M., & Schneider, A. (2020). Format matters: An introduction to format studies. In M. Jancovic, A. Volmar, & A. Schneider (Eds.), *Format matters: Standards, Practices, and politics in media cultures* (pp. 7–22). Meson Press. <https://doi.org/10.14619/1556>
- Voskuhl, A. (2004). Humans, machines, and conversations: An ethnographic study of the making of automatic speech recognition technologies. *Social Studies of Science*, 34(3), 393–421. <https://doi.org/10.1177/0306312704043576>
- Walker Rettberg, J. (2020). Situated data analysis: A new method for analysing encoded power relationships in social media platforms and apps. *Humanities & Social Sciences Communications*, 7(5), 1–13. <https://doi.org/10.1057/s41599-020-0495-3>
- Wellington, C. (2021, June 20). CVPR '21 WAD keynote – Carl Wellington, Aurora. *YouTube*. https://www.youtube.com/watch?v=YQ_yar_5Q98
- Wilmott, C. (2016). Small moments in Spatial Big Data: calculability, authority and interoperability in everyday mobile mapping. *Big Data & Society*, 3(2), 1–16. <https://doi.org/10.1177/2053951716661364>
- Wilmott, C. (2020). *Mobile mapping: Space, cartography and the digital*. Amsterdam University Press. <https://doi.org/10.2307/j.ctvx8b7zc>

5

Demonstrating Decisions: Waymo's World

In October 2020, Waymo lifted the lid on their autonomous vehicle operations in Phoenix, Arizona. Throughout 2019, and nine months of 2020, their vehicles had been involved in 47 so-called contact events, across 6.1 million miles with a trained operator present, and a further 65,000 miles (or roughly 3095 miles/month) without (Schwall et al., 2020, p. 1). These 47 incidents covered a range of different types of contact, but of these 47, only 18 actually happened. The remaining 29 were 'simulated events' predicted by Waymo's own counterfactual calculations. In these cases, trained operators had assumed control before an actual incident had occurred, thus preventing any subsequent contact event. In so doing, Waymo engineers would then run their own simulation(s) to determine whether a contact event *would* have happened. In 2023, Waymo lifted the lid again, in celebration of 1 million 'rider-only' (i.e. autonomous) miles, further evidencing the occurrence of contact events across their respective fleets in Phoenix and San Francisco (Victor et al., 2023).

In this chapter I want to explore the significance of these 'what if' scenarios and discuss how various kinds of calculations are not only central to testing the capabilities of autonomous vehicles but also, in their public release, central to *demonstrating* the safety of them too. In particular, I

want to focus on the design of algorithmic ‘path prediction’ processes meant to extrapolate intended road user trajectories. These simulated events constitute a novel kind of failure in the digital world: one that is paradoxically meant to document (present, probable) operational error and persuade of (future, possible) operational success.

In this, I argue that as these error events are only ever inputs for counterfactual calculations, assessment of the decision-making abilities of such vehicles is never complete: what Aradau and Blanke (2021) consider as the ‘optimization’ of error. In this, decisions are never ‘placed beyond doubt’ (Amoore, 2020, p. 137), only ever *deferred* and thus never actually, definitively, reached. In such a formulation, crashes are only ever rendered a partial, correctible state: the raw data for iterative ‘what if’ scenarios.

It is through this simulation work—through the deferral of decisions—that Waymo manages or ‘suspends’ the meaning of autonomous vehicle crashes. By nominally demonstrating the safety of their autonomous vehicles in such a way, Waymo is able to engineer, in a quite literal sense, internal and ultimately public knowledge and understanding of Waymo vehicle’s decision-making abilities. In effect, to be able to control that knowledge of their capabilities that is produced, circulated, and validated. This is, in short, what I call here ‘Waymo’s world’.

Whilst arguably a new context, computer simulations have always been used to manage operational outcomes. Put simply, simulations are always a management tool and a tool for management to make operational decisions. In short, that simulation runs and iterations are, I argue, always infused with discursive management work in which the meaning and significance of the simulation outputs are always modifiable and modified, according to organizational interests. In the case of Waymo, as this chapter contends, simulations—these ‘what if’ scenarios—are specifically used as a reputational resource, critical for demonstrating the safety of Waymo vehicles, and autonomous vehicles more generally.

Uber ATG: The Backstory

To start this story, it is necessary to return to 2018. As Elaine Herzberg was walking across Northbound Mill Avenue in Tempe, Arizona, she was hit and killed by a modified Volvo XC90 testing Uber ATG's developmental automated driving system (ADS), equipped with 20 ultrasonic sensors, 10 cameras, 8 radar sensors, and 1 lidar unit. The subsequent National Transportation Safety Board (NTSB) report into the crash, released in November 2019, revealed the harrowing details of the accident, that 'in the 5.6 seconds before Herzberg was hit, she was classified by the ADS on *ten separate occasions*, with each classification yielding a different possible *trajectory* [or path] Herzberg might take across the road' (Hind, 2022, p. 66, emphasis added). At no point did the ADS confidently recognize Herzberg. To continue:

On the first occasion, Herzberg was detected by the radar system as a *Vehicle*. 0.4 seconds later, she was detected by the lidar system and deemed to be a static object, putting her into the category of *Other*. One second later she is classified again as a *Vehicle*, but nonetheless is still presumed to be static. 2.6 seconds before impact, the ADS reclassifies her for a fourth time; this time as a *Bicycle*, deciding the bicycle by her side is being ridden. With 2.5 seconds left the system finally predicts she is moving, yet only through an adjacent lane to the test vehicle. 1.5 seconds before impact she is again classified as *Other*, and all previous trajectories are 'reset'. She is once again deemed to be a static object. At 1.2 seconds before impact she is reclassified for a final time, now as a *Bicycle*, with the ADS predicting she is in the direct path of the test vehicle. Now too late to safely execute an emergency avoidance strategy, the ADS initiates 'action suppression' designed merely to mitigate the effects of an impact. 0.2 seconds before Herzberg is hit, action suppression ends and the system issues an auditory warning. 0.02 seconds before impact, the vehicle operator (VO), Rafael Vasquez, takes control of the steering wheel; now powerless to prevent the fatal crash. (Hind, 2022, pp. 66–67)

Perhaps the most consequential part of the account above is that '1.5 seconds before impact [Herzberg] is again classified as *Other*, and all previous trajectories are "reset"' (Hind, 2022, p. 67). Thus, instead of the

ADS being able to incorporate what the NTSB report called user ‘tracking histories’—memory of the previous categorizations the ADS has made for the road user—into its calculations, each new categorization (as *Vehicle*, as *Bicycle*, as *Other*) generates an entirely new ‘path trajectory’ for the road user.

The significance of this is two-fold. Firstly, that the calculation of these path trajectories or ‘predictions’ constitutes the generation of what Waymo refers to as ‘what if’ scenarios, that is possible future states of, and interactions between, an autonomous vehicle and other road users. It is through the simulation of these ‘what if’ scenarios that the making of decisions is only ever deferred. But secondly, that in Uber ATG’s failure to design a system capable of incorporating user tracking histories, consequently leading to the death of a pedestrian, they also severely dented wider public confidence in the safety of autonomous vehicles, both in the US and across the globe. The result was a concerted attempt by Waymo (amongst others) to begin to *demonstrate* the safety of their own autonomous vehicles. In order to do so with full control over such a demonstration, Waymo turned to the world of computer simulation.

Simulation and the Management of Operations

The world of computer simulations can reasonably be said to have emerged through the discipline of operational research (OR), closely connected to cybernetics. OR, put simply, was the realization and deployment of cybernetic thinking within an operational context. As the *Journal of the Operational Research Society* (JORS) suggests, real-world applications of OR incorporate a whole breadth of areas of business and government, from forecasting and inventory management to project management and scheduling (JORS, 2022). The specific environments in which OR is typically used include energy, finance, manufacturing, and transportation (JORS, 2022). The specific technical approaches that are ordinarily grouped under the OR umbrella include ‘decision support systems, expert systems, heuristics, networks, mathematical programming,

multicriteria decision methods, problems structuring methods, queues, and simulation' (JORS, 2022, p. n.p.). In short, that OR is in part concerned with the development of both (a) computer simulations and (b) systems for supporting decision-making processes.

It is this antecedent work within OR, operationalizing simulation programming from the 1950s onwards, that can be brought to bear on contemporary simulation work carried out by Waymo. Using the first electronic stored-programme computers in the world, such as early models made by Ferranti, OR practitioners became simulation pioneers, developing prototypes for subsequent kinds of simulations beyond these original settings. This simulation work sought to offer an abstracted, exhaustive, environment in which different operational scenarios could be modelled and played out, designing many of the principles to be followed later by the likes of Waymo in simulating, and predicting, the paths taken by road users within a particular driving environment. What I want to argue here is that Waymo uses simulations to actively manage reactions and responses to their test operations, in ways that are somewhat, but not entirely, novel.

The original context for this early computerized simulation work was United Steel, a steelmaking company based in South Yorkshire, UK, that later became nationalized as British Steel Corporation in 1967. Ten years prior, in 1957, United Steel set up the Department of Operational Research and Cybernetics, arguably the first company department dedicated to cybernetics and OR in the world. Based at 'Cybor House' (standing for CYBernetics and Operational Research) in Sheffield, an interdisciplinary team of researchers, including 'three psychologists, an anthropologist, two zoologists, a philosopher and a classicist – as well as the range of scientific disciplines more normally (now) associated with an Operational Research department' (Hollocks, 2006, p. 19), began to explore the possibilities of applying new computational techniques to the steelmaking process. The new department was set up by Stafford Beer, who had convinced a hard-headed steelmaking firm to form a computational research and development (R&D) unit, composed of mystical minds.

In order to perform this experimental work, the unit acquired a 50% share in a Pegasus II computer, made by Manchester-based Ferranti,

ordered by the University of Sheffield (Hollocks, 2006). As Hollocks (2006) and Disley (1997) attest to, one of the most successful projects devised by the unit at Cybor House was a so-called General Simulation Program (GSP), designed to simulate every part of United Steel's steel-making process across sites in South Yorkshire and Lincolnshire. As Hollocks recounts, the General Manager of the Appleby Frodingham Works in the Lincolnshire town of Scunthorpe, George Elliott, in his own book, *Practical Ironmaking* (Elliot & Bond, 1959), had commented that 'in spite of the fashionable worship of such things as Operational Research, Automation, Cybernetics and the other catch words which be-devil industry, it is believed that the iron-works will be one of the last places where the practical man will be king' (quoted in Hollocks, 2006, p. 20). In other words, that despite managing to convince United Steel's senior management to fund Cybor House, many in the steelmaking industry were far from convinced that such computational research would be of actual, operational value.

Yet fairly quickly, United Steel became dependent upon the work of Cybor House and, specifically, the role of GSP in modelling United Steel's different steelmaking plants. In the first instance, the development of GSP required the establishment of an operational ontology to refer to different parts of the steelmaking process. As the first version of the handbook supposed, GSP was designed to simulate 'events' where 'there was a discrete change of activity' (Tocher et al., 1959, p. 6) within the steelmaking process. Generating these events, comprised of discrete changes of activity, were a 'collection of machines' each 'capable of taking one of a set of clearly defined states' (Tocher et al., 1959, p. 6). Each 'cycle' within the steelmaking process (e.g. the Bessemer cycle, casting pit cycle) thus comprises activities that are time-dependent (e.g. the pouring of molten steel) and a selection of 'stores' (e.g. casts) without a temporal dimension. Activities, in the words of the GSP handbook authors, were thus simply just 'groups of machines in certain specified states' (Tocher et al., 1959, p. 7). Connecting each stage of this overall process together, of course, was a sustained flow of metal, starting with the blast furnace iron and finishing with steel ingots (see Fig. 5.1). Put together, machines, events, activities, states, cycles, stores, and flows could be said to comprise the entirety of the steelmaking process to be modelled and thus simulated.

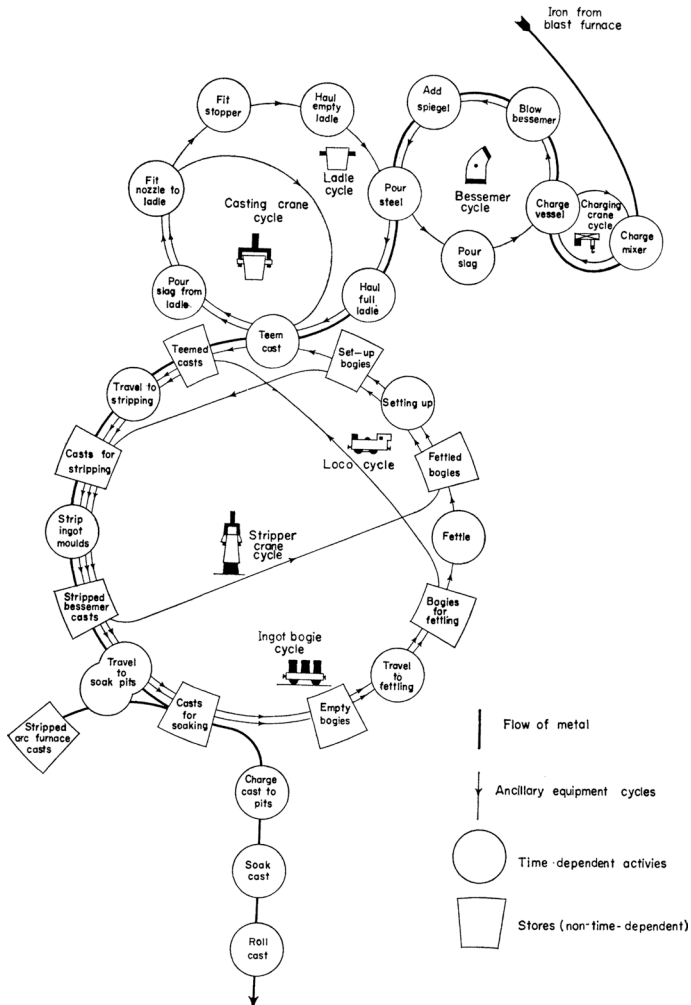


FIG. 4. Simplified flow diagram of activities. Acid Bessemer steel-making plant.

Fig. 5.1 A simplified flow diagram of activities at an Acid Bessemer steel-making plant. (Source: Tocher, 1960)

This operational ontology can be understood as an example of what Phil Agre later calls ‘grammars of action’ (Agre, 1994). In other words, that the researchers at Cybor House were engaged in ‘the practice of constructing systematic representations of organizational activities’ (Agre, 1994, p. 108). As Agre intimates, efforts to map such activities were not necessarily new to the computer age but were clear descendants of time and motion studies devised for industrial purposes at the beginning of the twentieth century (McKinlay & Wilson, 2012). Despite, or really because of, this lineage, the computer became the ideal technology for employing ‘formal “languages” for representing human activities’ (Agre, 1994, p. 108), such as the manufacture of steel. Establishing a ‘grammar’ of activities consisting of ‘minimum replicable units’ (Agre, 1994, p. 108)—say events, states, stores, and flows—would thus help to construct such systematic representations. With or without computers, such work would help ‘bring to management’s notice’ any outstanding ‘redundancies [or] other inefficient patterns of activity’ (Agre, 1994, p. 108). Yet, specifically *with* them, managerial decision-making could—in theory, at least—take place with greater speed and efficiency.

It is not the purpose of this section to restate the principles of Agre’s (Agre, 1994) ‘capture model’ itself to which Agre argues grammars of action typically serve. Nonetheless, it is worth mentioning that the five stages of the model (analysis, articulation, imposition, instrumentation, and elaboration) can be said to have been readily applied to United Steel operations. Moreover, that GSP was principally responsible for enabling the ‘normative force’ (Agre, 1994, p. 110) of the operational ontology developed by the Cybor House team to map the company’s steelmaking operations. Through this initial mapping exercise, otherwise the building of a steelmaking ‘ontology’, United Steel managers were able to rationalize and streamline the steelmaking process across their many sites and steelplant types. In simulating the steelmaking process, United Steel managers were better placed to make executive decisions regarding the objectives of the firm itself, including decisions to optimize the running of the interconnected ‘machines’ within the steelmaking process or install new steelplant systems altogether.

Path Predictions: Actualizing the Virtual

One fundamental difference between the simulation work at United Steel and Waymo is the difference between the modelling of 'discrete-event' processes and 'visual-interactive' approaches (Steinhoff & Hind, 2024). Whilst steelmaking takes place in a linear, sequential fashion (albeit with interlocking cycles), the movement—and therefore the simulation—of an autonomous vehicle must incorporate near-infinite forking paths and interactions. Considered as early as the 1930s, this 'field of safe travel' any vehicle must proceed through 'consists at any given moment, of the field of possible paths which the car may take unimpeded' (Gibson & Crooks, 1938, p. 454). As JJ Gibson and Crooks write, 'we may assume that driving is a type of locomotion through a "terrain" or field of space' (Gibson & Crooks, 1938, p. 454) and that 'locomotion is therefore guided chiefly by vision' with 'this guidance ... given in terms of a "path" within the visual field of the individual such that obstacles are avoided and the destination ultimately reached' (Gibson & Crooks, 1938, p. 454). Through what amounts to an operational ontology or grammar of action, 'these concepts of *terrain*, *destination*, *obstacle*, *collision* and *path* should be applicable to any type of locomotion ... [including] the operator of an automobile' (Gibson & Crooks, 1938, p. 454, authors' emphasis).

Central to simulating these forking paths and interactions—pivotal to 'all algorithmic arrangements' (Amoore, 2020, p. 9)—are so-called path trajectories or path predictions. In short, algorithmic calculations that predict the future state, and location, of objects or road users in a wider driving environment. Most autonomous vehicle systems need to generate path predictions, and all that do, do so with the assistance of an array of sensors. Usually this is some arrangement of high-definition cameras, radar, and lidar in order to cover short-, mid-, and long-range distances and avoid any blind spots. Once sensor data is captured it tends to be processed by a 'perception' module in on-board software, which 'detects and tracks individual actors and objects in order to generate estimates of their position, orientation, and velocity and register other attributes that may inform their future motion' (Uber ATG, 2020, p. 29). Path predictions usually appear in a 'prediction' module, 'which applies different

models of behavior for different actor and object *classes*' (Uber ATG, 2020, p. 29, emphasis added). Quoting again from Uber ATG's last Safety Report in 2020,¹ before being sold to Aurora (Korosec, 2020):

The Prediction software considers and presents multiple anticipated motion paths for objects — i.e. possibilities of what the tracked 'object' might do next — to the Motion Planning software, including intents that the system predicts may put the actors or objects in the self-driving vehicle's path, even when the self-driving vehicle has the right-of-way. The Prediction system seeks to determine the probabilities of multiple future paths for each actor in the scene. The Motion Planning system then uses these probabilities to effect an appropriate amount of caution in response to less predictable actors or objects. The system performs these predictions many times a second so as actors change direction or intent the system continually reassesses their likely next move. (Uber ATG, 2020, p. 29).

A good example of what these predictions look like visually can be seen in work by Christian Pek and colleagues (Pek et al., 2020), on the use of formal verification techniques to guarantee the legal safety of autonomous vehicles when executing decisions. What is interesting here is how they bear a great resemblance to diagrams of the 'fields of safe travel' drawn by Gibson and Crooks (1938). In both cases, the authors establish an operational ontology/grammar of action in order to facilitate a prediction of the future—safe—path of the vehicle (Fig. 5.2).

From an ethical perspective, rather than a strictly technical one, path prediction work might be mistaken for the infamous 'trolley problem', where a 'consequentialist' ethics engenders a binary approach to decision-making (Ganesh, 2017). Proceeding instead in a probabilistic fashion—where paths are weighted according to possible risk—this work instead typifies the 'ontologically multivalent' (Ganesh, 2022, p. 194) nature of path simulation. Pek et al. (2020) and others work on path predictions demonstrate this work in action.

¹ A blog post announcing the publication of the report from the Uber ATG Safety Team can be found here: <https://medium.com/@UberATG/uber-atg-releases-2020-safety-report-575db33f2bd7>




FIG. 1. THE FIELD OF SAFE TRAVEL AND THE MINIMUM STOPPING ZONE OF A DRIVER IN TRAFFIC

(If, in this and the following figures, the page is turned around and the figure is viewed from what is now the right, the reader may the better be able to empathize the situation, since he will then have the point of view of the driver of the car whose field of safe travel is under discussion.)

Fig. 5.2 A visual depiction of the 'field of safe travel'. (Source: Gibson & Crooks, 1938)

In the scenario of a 'jaywalking' pedestrian (i.e. the same scenario as in Tempe, Arizona), the autonomous vehicle is depicted on the left of the image, with its own 'intended' trajectory projected through a single lane from left to right in black. Other road users/objects are present, including two trucks (larger blue rectangles), two pedestrians (blue dots) either side of the road, and one other (moving) vehicle in the other lane ahead (small blue rectangle). In this 'fail-safe' approach, predicted paths are calculated, generating predicted 'occupancy sets', larger areas where the other road users might plausibly be over the course of 7.8 seconds. With the high level of safety they employ in their modelling—arguably a higher level than present in the Uber ATG crash in 2018—any of the users in this scenario could end up occupying the corresponding areas in blue.

In particular, the pedestrian crossing the road (ID 323 in Fig. 5.3) could occupy any point within the broad blue area extending from the right-hand side: on the pavement, in the same lane as the vehicle, in the parallel lane, next to one of the parked trucks, etc. Opting for a fail-safe approach, the vehicle can only proceed as far as the red line takes it, before encroaching the occupancy set of the jaywalking pedestrian. As they aptly remark:

Even though it is illegal for pedestrians to jaywalk [in some countries, i.e. USA and Germany], that is, to cross the road in the presence of traffic, pedestrians are occasionally inattentive and cross directly in front of passing vehicles. If the prediction of the autonomous vehicle does not include this behaviour, a fatal accident could occur. (Pek et al., 2020, p. 522)

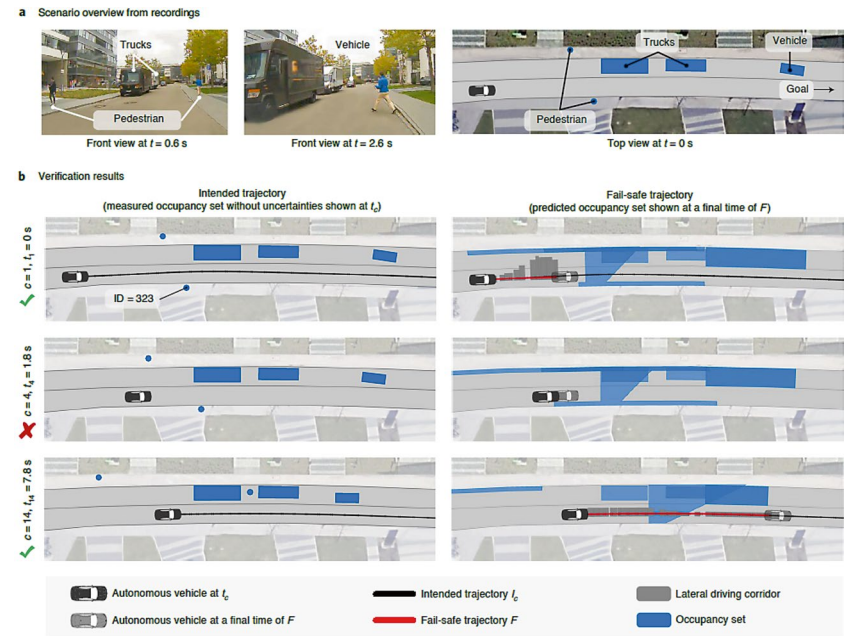


Fig. 4 | Results of Scenario II (jaywalking pedestrian). a, Camera images and top view of the scenario. b, Verification results of selected verification cycles c. The intended trajectory t_c is only shown if it is successfully verified. Credit: Google, GeoBasis-DE/BKG (satellite images).

Fig. 5.3 A visual depiction of a path trajectory taking account of a 'jaywalking' pedestrian. (Source: Pek et al., 2020)

Thus, path predictions are not only critical to ensuring such incidents do not occur but also a particular part of the decision-making process through which potentially infinite outputs can be modelled. Here, the generation of safe path trajectories is conditional both on the validity of sensor data being fed through the operational 'pipeline' and the latitude with which occupancy sets are generated. Restricting something like the future 'plausibility range' an 'errant' pedestrian might wander into ordinarily produces an autonomous vehicle less equipped to deal with more unpredictable scenarios or less predictable road users. In other words, building an autonomous vehicle for a more rigid social environment, where other road users are liable to be seen to 'err' from established expectations than in other scenarios or situations. Or, indeed, in situations or scenarios populated purely with human drivers, perhaps better equipped to deal with 'errant' behaviours. In this modelling work, questions of risk and safety are necessarily contingent on a range of possible factors from vehicle speed to user classification. In the following I will discuss the modelling conducted by Waymo and their specific delineation of so-called contact events.

Waymo Case: 2019–2020

Contact events are Waymo's way of referring to accidents involving their autonomous vehicles, using a standard severity scale derived from the ISO on road vehicle safety (ISO 26262).² The term itself engenders a discursive shift away from directly attributing blame, liability, and responsibility, as well as immediately negative connotations associated with terms like crash, accident, and collision. In using the term contact event, Waymo already establishes a certain neutrality of relations between its vehicles and other affected parties. Used akin to something in particle physics, 'contact' between Waymo vehicles and other road users only seems to occur if two directionless, motivation-less, objects happen to come together.

²From S0, 'no injury expected', to S3, 'possible critical injuries expected'.

Waymo divides these contact events into further categories: whether constituting ‘single vehicle events’ (i.e. the Waymo vehicle only) or ‘multiple vehicle events’ (i.e. involving the Waymo vehicle and other non-Waymo vehicles). Different types of collision are also recorded, from those related to single vehicle events such as involving a ‘road departure’ (i.e. the Waymo vehicle leaves the road altogether), ‘striking a pedestrian/cyclist’ or being ‘struck by pedestrian/cyclist’. Five different contact events are listed for multiple vehicle incidents including ‘reversing’ incidents, ‘sideswipes’, ‘head-on and opposite direction sideswipes’, ‘rear end’ incidents (where another vehicle hits the back of the Waymo vehicle), and ‘angled’ incidents (e.g. if another vehicle hits the Waymo vehicle coming out of a side street, etc.) (Schwall et al., 2020). As Taina Bucher considers, categorization work functions as a ‘powerful mechanism’, not only ‘in making data algorithm-ready’ (Bucher, 2018, p. 5) but consequently in making engineers, and other interested parties, ready and able to interpret the actions of autonomous agents.

In the first ever public report on incidents involving Waymo vehicles, over 21 months from 2019 and the first nine months of 2020, 47 contact events were recorded (Schwall et al., 2020). Most were multiple vehicle events, according to Waymo’s classification, with the majority involving other vehicles either sideswiping the principle Waymo vehicle or hitting one from behind. However, in the events documented in the report, only 18 of these 47 contact events actually happened. A subsequent 29 were *simulated* contact events that Waymo believes would have happened if a human operator did not assume control at the time. Out of these 29 simulated events, 9 were sideswipes and 14 were angled incidents. No incidents at levels S2 or S3, the most critical levels on the severity scale were recorded, either actually or simulated, and a majority (30/47) were at level S0, that is that ‘no injury [was] expected’ (Fig. 5.4).

Counterfactuals as Demonstration Devices

The distinction between actual and simulated contact events is significant. The latter are the result of what Waymo researchers refer to as ‘counterfactual (“what if”) simulations’ (Schwall et al., 2020, p. 1). It is not a

Row#	Event type	Manner of Collision ("Other" = non-Waymo vehicle)	Waymo-involved collision-relevant contacts by ISO 26262 severity classification Actual & simulated event counts (Totals in Bold)					Human Crash Statistics (Non-Waymo Data)	
			S0	S1 (no airbag deployment)	S1 (airbag deployment any vehicle)	S2	S3	Collision % Contribution US*	Fatal Collision % Contribution U.S. (Maricopa Cnty, AZ) **
1	Single Vehicle Events	Road Departure, Fixed object, Rollover	0	0	0	0	0	20%	27% (21%)
2		Striking a pedestrian/cyclist	0	0	0	0	0	2%	33% (41%)
3		Struck by pedestrian/cyclist	1 (actual) 2 (sim)	0	0	0	0	<0.5%	1% (1%)
4	Multiple Vehicle Events	Reversing	1 (actual) 1 (sim)	0	0	0	0	1%	<0.1%
5		Other reversing, Waymo straight	1 (actual) 1 (sim)	0	0	0	0		
6		Waymo reversing, Other straight	0	0	0	0	0		
7		Sideswipe (Same Direction)	1 (actual) 8 (sim)	1 (sim)	0	0	0	11%	1% (1%)
8		Other lane change, Waymo straight	1 (actual) 7 (sim)	0	0	0	0		
9		Waymo lane change, Other straight	1 (sim)	1 (sim)	0	0	0		
10		Head-on + Opposite Direction Sideswipe	0	0	1 (sim)	0	0	5%	9% (7%)
11		Rear End	11 (actual) 1 (sim)	1 (actual) 1 (sim)	2 (actual)	0	0	34%	5% (5%)
12		Other striking, Waymo struck (stopped)	8 (actual)	0	0	0	0		
13		Other striking, Waymo struck (slower)	2 (actual)	1 (actual)	1 (actual)	0	0		
14		Other striking, Waymo struck (decelerating)	1 (actual)	1 (sim)	1 (actual)	0	0		
15		Waymo striking, Other struck (stopped)	0	0	0	0	0		
16		Waymo striking, Other struck (slower)	0	0	0	0	0		
17		Waymo striking, Other struck (decelerating)	1 (sim)	0	0	0	0		
18		Angled	4 (sim)	6 (sim)	1 (actual) 4 (sim)	0	0	27%	24% (24%)
19		Same direction - Other turns across Waymo straight travel	0	2 (sim)	0	0	0		
20		Same direction - Other turns into Waymo straight travel	3 (sim)	0	2 (sim)	0	0		
21		Opposite direction - Other turns across Waymo straight travel	0	0	1 (sim)	0	0		
22		Opposite direction - Other turns into Waymo straight travel	0	0	1 (sim)	0	0		
23		Straight crossing paths	0	1 (sim)	1 (actual)	0	0		
24		Same direction - Waymo turns across other straight travel	1 (sim)	3 (sim)	0	0	0		
25		Same direction - Waymo turns into other straight travel	0	0	0	0	0		
26		Opposite direction - Waymo turns across other straight travel	0	0	0	0	0		
27		Opposite direction - Waymo turns into other straight travel	0	0	0	0	0		
28		Total	14 (actual) 16 (sim)	1 (actual) 8 (sim)	3 (actual) 5 (sim)	0	0	100%	100% (100%)

Fig. 5.4 Classification of Waymo-involved collisions. (Source: Schwall et al., 2020)

method they suggest is the norm within autonomous vehicle testing but contend is a method ‘increasingly used’ (Schwall et al., 2020, p. 2) within this context, in order to ‘provide an opportunity to study what would likely have occurred in a specific scenario had the Waymo Driver [i.e. the vehicle] remained engaged’ (Schwall et al., 2020, p. 2).

Generating these ‘what if’ scenarios, according to Waymo, is ‘significantly more realistic’ (Hawkins, 2020, p. n.p.) than wholly ‘synthetic’ alternatives (i.e. those software-generated without a real-world trigger), and of course bear none of the immediate danger of actual crashes. Consequently, they are designed to constitute a central safety feature of Waymo’s autonomous vehicle testing programme.

Simulating, classifying, and publicizing these contact events are an attempt to *demonstrate* the safety of Waymo’s autonomous vehicle operations. In their own words, they write that ‘[t]he goal of this transparency is to *contribute to broad learning* with the industry, policymakers, and the

public; promote awareness and discussions; and foster greater public confidence in automated vehicles' (Schwall et al., 2020, p. 11, emphasis added). The demonstration of the functionality of a specific technology here is nothing new, with a long history in modern science and technology, from air pumps (Shapin & Schaffer, 1985) to vaccines (Latour, 1988) to AI (Hong, 2022).

A large part of the need for such demonstration work is because of the contingency of the operations being simulated. Here, the operational ontology required for the simulation of the (open) path of a vehicle differs from the (closed) flow of molten blast furnace iron. Whilst there may be comparable direct risk to both operators (of autonomous vehicles and steel plants) in terms of operational inefficiencies or poor executive decision-making, only Waymo must publicly demonstrate the safety of their operations—hence the work being justified by appointed Waymo engineers.

The staging of contact events entails what Noortje Marres (2020, p. 8) has referred to as a 'distinctive mode of publicity' in which 'contingent and contextual occurrences and encounters' are deployed to reassure interested parties not simply that autonomous vehicles are nominally safe but that they are safer than the autonomous vehicles elsewhere—like those scrutinized in the NTSB report into the Uber ATG crash in March 2018. As Aradau and Blanke (2021, p. 2) help us to understand, contact events as particular kinds of automotive mistakes or errors are 'not limited to laboratories or scientists' debates, but [are] invoked in arguments that shape public debates', in this case about the safety of autonomous vehicles. The collection of this data and the running of counterfactual calculations can be considered a response by Waymo to the reputational damage done by Uber ATG in the year previous: data documented in the report begins on January 1, 2019, 11 months after the crash. Hence, one can consider this work as a *re-staging* or a *re-situating* of the public debate around the safety of autonomous vehicles.

In the released data, there are no incidents that either appear in more severe categories (involving possible critical injuries or death—as in the Uber ATG case) nor are there any recorded contact events involving the Waymo vehicle striking a pedestrian or cyclist (again, as in the Uber ATG case). This, I contend, is not coincidental but the result of Waymo

responding very specifically to the public debate that emerged in the aftermath of the 2018 crash, and the specific fear that pedestrians and/or cyclists would be most at risk from autonomous vehicles. Tellingly, the Waymo data suggests that other human drivers are the greatest such risk, with Waymo’s stated mission being ‘to reduce traffic injuries and fatalities’ as well as ‘improve mobility for all’ (Schwall et al., 2020, p. 1).

Waymo’s addition of ‘human crash statistics’ to their own contact event data offers a favourable comparison between the (safer) Waymo Driver and other (more dangerous) human drivers. For instance, in how the registering of zero contact events in the two first categories of single vehicle events (‘road departure’ and ‘striking a pedestrian/cyclist’) are said to ‘combine to contribute approximately 60% of all human-driven fatal collisions on sub 45mph urban roadways both nationally and within the Maricopa County, Arizona, where Waymo’s ODD [operational design domain] is located’ (Schwall et al., 2020, p. 6) (Fig. 5.5).

These different elements are representative of what Gregg Culver (2018, p. 145) has called the ‘profound escapism’ that autonomous vehicles will ‘usher in an era of fatality-free automobility’. In contrast to this chimera of the ‘fatality-free automobile’, Culver frames cars (autonomous

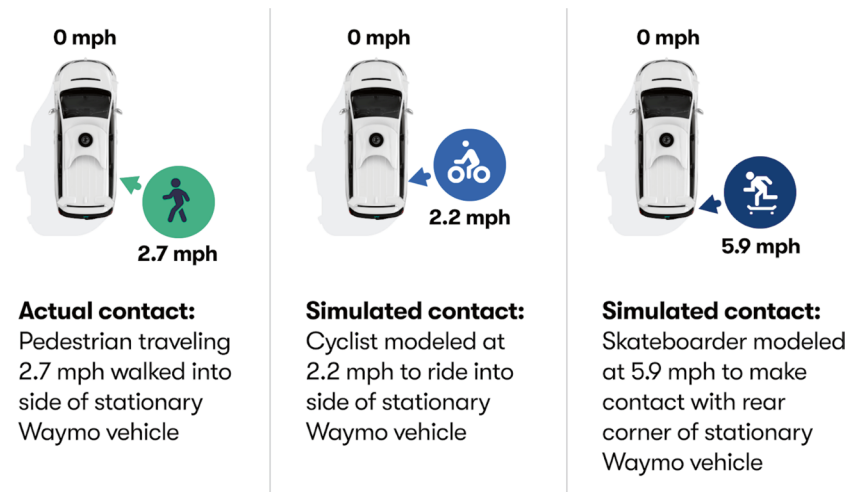


Fig. 5.5 Actual and simulated contact events involving pedestrians or cyclists. (Source: Schwall et al., 2020)

or otherwise) as offering forms of ‘vehicular violence’ (Culver, 2018, p. 149):

In the calculus of mobility violence, vulnerable users and motorists present strikingly different packages of threats and vulnerabilities, such that a collision between a motorist and a vulnerable road user can often mean death to the latter while leaving the former physically unharmed. (Culver, 2018, p. 149)

‘Unpredictable’ pedestrians figure so prominently in autonomous vehicle research because they literally stand in the way of autonomous vehicle success and the fantasy of ‘fatality-free automobility’.

On an operational level, the generation of these simulated ‘what if’ events is a way of *learning* through the kinds of failures in which likely errors are anticipated, extrapolated from current states, averted, and subsequently continued within a virtual environment. As Aradau and Blanke (2021, p. 1) write: errors are ‘now inherent to vernacular modes of knowledge and mundane practices of human-machine interaction’. The simulations involve a number of stages after the ‘sliding doors’ moment between the actual trajectory *not taken* by the Waymo Driver after the point at which a human operator is alerted (and the Waymo Driver ‘dis-engages’) and the *simulated trajectory* (or ‘post-disengagement simulation’ [Schwall et al., 2020, p. 3]) subsequently generated.

I understand these events as ‘quasi-virtual’ because of their real-world triggers, using actually existing agents. In this, they perhaps complicate the discussion Marres (2020, p. 5) has regarding Karin Knorr-Cetina’s (2009) idea of ‘synthetic situations’. As Schwall et al. (2020, p. 2, emphasis added) write: ‘[c]ounterfactual disengagement simulations can be significantly more realistic than simulations that are created *entirely synthetically* because they use the actual behavior of the autonomous vehicle and other agents up to the point of disengagement’. In this, ‘actual behaviour’ is used as a baseline or a ‘data ground truth’ (Amoore, 2020, p. 137), for counterfactual calculation. While Schwall et al. (2020) understand synthetic as ‘non-actual’, Knorr-Cetina (2009, p. 66) considers it as a kind of augmented, mediated, reality enabled by ‘scopic’ systems (Knorr-Cetina, 2009, p. 64).

Here a difference may rest on the idea of 'entire' synthesis, as opposed to 'partial' synthesis. Understanding them as quasi-virtual avoids such a complication altogether, drawing instead on the Deleuzian notion of virtuality as a trajectory or path towards actualization (Deleuze, 1981; Shields, 2002). Sprenger (2020), in a similar vein, considers such work as dependent on generating 'virtual probabilities' (Sprenger, 2020, p. 619), as the 'microdecisions that underpin autonomous technologies are both an element and effect of the ... virtualization of an environment into probabilistic models' (Sprenger, 2020, p. 621).

Generating these simulations requires a number of sequential steps. The first involves simulating vehicle *motion* (offline, on the same software as present in-vehicle) in which the vehicle's 'pre-disengage position, attitude, velocity, and acceleration along with the autonomous vehicle's recorded sensor observations' (Schwall et al., 2020, p. 3) are used. For the data used in the report itself, these simulations were carried out on a 2019 version of the software, rather than any current version (avoiding, therefore, using a version that had already benefited from learning from these incidents in the first place).

After this motion data is simulated, 'a check is performed to determine if the simulated positions of the autonomous vehicle overlap at any point with the recorded positions of other agents' (Schwall et al., 2020, p. 4). In effect, Waymo is implementing a version of occupancy calculations detailed by Pek et al. (2020), ensuring that the simulated scenario would not have resulted in a 'potential collision' (Schwall et al., 2020, p. 4). Although, again, 'potential' is important to highlight here because of the various possible path trajectories and positions these other simulated agents might have taken, and been in. What is interesting here is that Waymo uses a jaywalking scenario to illustrate this process (akin to the Uber ATG crash, and just like Pek et al., 2020 discuss). In their vulnerability, jaywalkers are decidedly powerful.

At this point the question of modelling other users becomes important, because of how 'the counterfactual behaviour of the autonomous vehicle may have elicited a different response from other nearby roadway users' (Schwall et al., 2020, p. 4). As the authors write, the modelling of other user responses is difficult, not least because of the range of possible responses (and reaction times) drivers have to specific stimuli,

such as a vehicle pulling out of a junction. To deal with this difficulty the authors have resorted to *simplicity*: ‘the results reported in this paper are based on deterministic models that generate a single response to a given input’ (Schwall et al., 2020, p. 4). That is, they forgo what they call ‘probabilistic counterfactual outcomes’ (Schwall et al., 2020, p. 4) that better account for the range of reactions, because of the complexity of doing so. This narrowing naturally constrains the number of possible contact events that can be generated from any one input, ensuring the Waymo Driver appears safer than it might be. This ‘decisional determinism’ (Sprenger, 2021, p. 170) limits the range of possible outcomes, considered more favourable than any approach that generates multiple, possible results. The more possible outcomes, the greater possible risk for Waymo.

Waymo Case: 2020–2023

Three years later, after a purported one million miles of ‘rider-only’ (RO) operations, Waymo released another safety update (Victor et al., 2023). Twenty contact events were recorded across a period from September 2020—the final month covered in the first release—to January 2023, the month in which the company hit their one million milestone of ‘rider-only’ (autonomous) operations (Fig. 5.6). The data is derived from Waymo’s two operational areas (Phoenix, Arizona, and San Francisco) and two operational platforms (Chrysler Pacifica and Jaguar I-Pace) (Victor et al., 2023).

The report does not detail any counterfactual simulation work undertaken during this period—unlike in the original. As a result, all 20 contact events are actual incidents that occurred. The yearly breakdown of these actual contact events is: 1 in 2020 (September to December, only), 6 in 2021, 11 in 2022, and 2 in 2023 (January only). Twelve contact events involved the fourth generation of Waymo vehicles operating in Phoenix, six involved the fifth generation in Phoenix, and two involved the fifth generation of Waymo vehicles operating in San Francisco (Victor et al., 2023).

In January 2023, Waymo reached 1 million rider-only miles



Fig. 5.6 An infographic of Waymo reaching one million 'rider-only' miles. (Source: Waymo., 2023)

Using a calculation of injury risk referred to as the probability of maximum Abbreviated Injury Scale of 2 or greater ($p(\text{MAIS}2+)$), the contact events are assessed for their severity. As the report suggests, 'examples of AIS2 level injuries are concussions with no or brief loss of conscience, fractures to the sternum, and 2 or few rib fractures' (Victor et al., 2023, p. 8). Thus, the $p(\text{MAIS}2+)$ percentage likelihood for each contact event represents the chance of each incident resulting in such injuries to occupants within the vehicle. Only one such incident recorded a figure above 2%—the first contact event recorded in September 2020 in Phoenix. A further eight contact events recorded a figure between 1% and 2%, and a final 11 contact events were between 0% and 1%. Much like other aspects of the new report, as I will detail below, the inclusion of this injury scale can be considered an enhancement of efforts in the 2020 report (Schwall et al., 2020) to classify collisions according to an approved ISO severity scale (ISO 26262), used by the US National Highway Traffic Safety Administration (NHTSA).

Categorizing Conflicts, Apportioning Blame

Contact event risk is strictly based on a calculation of possibility of injuries to vehicle occupants. Arguably, this is not the aspect of such operations that is particularly in dispute. Indeed, the automotive industry at large—precisely the argument here—largely has no problem with maximizing the safety of the occupants of their own vehicles. The problem—one raised in both the Uber ATG case and the case of Cruise in San Francisco (see Chap. 8)—is that the minimal risk likelihood is not, and indeed *never*, extended to other, ordinarily more vulnerable, road users.

This is arguably another example of autonomous vehicle firms such as Waymo attempting, and failing, to demonstrate their vehicles *bear the weight of their own decision-making*.

The question that must also be asked—in lieu of mention of any simulated contact events—is whether the idea of the deferral of decision-making (the ‘non-decision’) is still applicable. As I detailed previously, the deferral of decision-making is about *suspending* and ultimately desiring to *manage* the otherwise doubtful meaning of the autonomous vehicle crash. The generation of simulated contact events (and the subsequent reporting of such) is one technique—but not the only technique—for doing so. Here this discursive management work might involve different techniques to achieve these end discursive goals, in which the meaning of such crash events is managed, suspended, and deferred. Indeed that, principally, the work being undertaken here is concerned namely with *demonstrating* the safety of particular, different kinds of autonomous vehicles, in different settings, under different conditions.

One of the additional ways in which this work is carried out, or at least reported on, in the 2023 report, is the development of what the authors refer to as a ‘conflict typology’ (Victor et al., 2023). Building on their 2020 effort, Waymo now details 16 such conflict ‘types’, where a conflict typology ‘describes the conflict partners, role (initiator or responder), and perspectives of each actor involved in a conflict’ (Victor et al., 2023, p. 28). In this, ‘conflict’ can be understood as the underlying basis for possible, future contact events, with any such actual contact event able to be categorized according to 1 of these 16 conflict types. Each conflict

type therefore concerns the state and trajectory of either the principle ('ego') vehicle in any given situation and/or other implicated road users such as pedestrians (Fig. 5.7).

These include previously defined types like 'Single Vehicle (SV)', now described as 'all actions (or lack thereof) where the ego vehicle is traveling in a trafficway' subsequently experiencing 'an in-trafficway interaction without a conflict partner (e.g. a rollover event) or an off-trafficway interaction (e.g. a road departure)' (Victor et al., 2023, p. 28), and 'Intersection Turn Into Path (ITIP)' (previously only referred to as an 'angled' collision), involving 'interactions that occur as a result of one of the actors







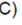



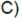

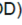



Conflict Group	Picture	Description
Single Vehicle (SV) 		Includes all actions (or lack thereof) where the ego vehicle is traveling in a trafficway but then experiences an in-trafficway interaction without a conflict partner (e.g., a rollover event) or an off-trafficway interaction (e.g., a road departure).
Front-to-Rear (F2R) 		Involves one road user interacting with another road user in the same direction and same travel lane.
Same-Direction Lateral Incursion (SDLI) 		Occurs when two roadway actors are traveling in the same trafficway but in initially different travel lanes at the time of the initial interaction due to lateral incursion by some actor.
Same-Direction Prior Circumstances (SDPC) 		Involves two roadway actors operating on the same trafficway in the same direction when one road user performs a lateral evasive action, experiences loss of control, or is involved in a prior collision that results in an interaction with the other road user.
Opposite Direction Lateral Incursion (ODLI) 		Occurs when a non-turning actor operating in the trafficway's intended travel direction interacts with another actor that is operating opposite of the travel direction in the same trafficway.
Opposite Direction Prior Circumstances (ODPC) 		Involves two roadway actors traveling in opposite direction trafficways in their respective trafficway's direction of travel when one road user performs a lateral evasive action, experiences loss of control, or is involved in a prior collision that results in an interaction with the other road user.
Turn into Path Opposite Direction (TIPOD) 		Occurs as a result of one actor changing vehicle-operated trafficways via a turning maneuver and interacting with another actor, where one of these actors is operating in the opposite direction of the trafficway's direction of travel.
Intersection Cross Traffic (ICT) 		Involves interactions that occur as a result of both actors changing or crossing over trafficways, and where the two actors cross paths with one another.

Fig. 5.7 Waymo's 'conflict typology'. (Source: Victor et al., 2023)

moving on to a trafficway via a turning maneuver into the path of another actor that is operating in the trafficway being turned into' (Victor et al., 2023, p. 28). In these and all other cases, there are varying degrees of action, interaction, and lack of action involved that variously implicate both the ego vehicle and other road users.

What this typology of 'conflict groups' serves to do is to construct a comprehensive operational ontology or grammar of action (Agre, 1994) through which path trajectories can be simulated. Whilst Waymo documented their conflict classification work in the 2020 report (Schwall et al., 2020), this was clearly a less sophisticated effort, with fewer conflict categories, and no accompanying typological figures. In developing this more sophisticated grammar, Waymo seeks to *suspend* and *manage* the outcome of both real and synthetic contact events, offering justifiable, plausible, arguments for why other vehicles—rather than the Waymo vehicle—were to blame for each such event.

The Waymo engineers subsequently discuss patterns of contact events discoverable in the data. Whilst these are not mutually exclusive, as they suggest, four conflict groups were found to be more common than the others: (Other Agent) Backing (BACK), Front to Rear (F2R), Contact with Objects in the Roadway (SV), and Opposite Direction Lateral Incursion (ODLI).

Other Agent Backing involves other vehicles backing into, or hitting, a Waymo vehicle. Eight out of twenty contact events in the data pertain to this conflict group, where 'in all backing contacts the Waymo vehicle was stationary at the time of impact' (Victor et al., 2023, p. 14). In these cases, it's inferred that the other drivers were responsible for the accident rather than the Waymo vehicle. In Front to Rear incidents, other vehicles were responsible for striking the Waymo vehicle from behind. Six contact events were considered F2R incidents, with the Waymo vehicle in each case 'either stationary or moving slowly at the time of impact' (Victor et al., 2023, p. 14). Once again, it's assumed the Waymo vehicle was blameless.

Interestingly, a further five incidents involved the Waymo vehicle hitting an object. All these incidents were regarded by Waymo as involving only 'non-fixed objects with low mass' (Victor et al., 2023, p. 14), including 'a construction pylon, shopping cart, swinging gate, plastic folding

sign, and an unoccupied miniature motorcycle' (Victor et al., 2023, p. 14). This rather humorous assortment of items comprises some of the items in the infinite 'prop stash' (Hind, 2019, p. 412) such vehicles have had—but rarely been able—to contend with, as other chapters consider.

Whilst the final common category, Opposite Direction Lateral Incursion, might be considered serious—essentially when another road user crashes into the vehicle from the opposite direction—this particular incident only involved a garbage truck driven at low speed. Occurring whilst 'negotiating a narrow passageway' (Victor et al., 2023, p. 15), the report makes explicit that despite being in the ODLI category, it bears little resemblance to 'common ... police-reported collisions that feature drivers drifting over the lane line while travelling at speed' (Victor et al., 2023, p. 15).

Waymo has gone to great lengths to apportion blame to all other road users for the noted contact events. The desired effect is to convey to the public that Waymo vehicles are safe, responsible, and never the cause of ordinary road accidents. By comparison, as the data is designed to demonstrate, human road users are ordinarily to blame for a suite of incidents, from harmless bumps to airbag-deployed crashes. The data is also designed to show a broad direction of travel, from the 4% p(MAIS2+) registered in the first incident in September 2020 to the 0% p(MAIS2+) figure in the last incident in February 2023.

How might we summarize these insights? Put simply: only other people, and other objects, are to blame. Whilst Waymo seems to understand this as a way to make everyone (and everything) else responsible for incidents their vehicles are involved in, it also reveals a fundamental truth of autonomous vehicles: the struggle to deal with the 'swarming social reality' (Hind, 2019, p. 412) of the road is real. The detailed typology of conflict types is intended to aid in the practical analysis of such incidents, enabling a more precise model for apportioning responsibility, of managing the response to contact events that have happened. The unintended effect of such is to offer even more granular understandings of how unprepared Waymo vehicles are for navigating complex environments—from passenger vehicles backing out of parking spots, to non-fixed, low mass objects like swinging gates and unoccupied miniature motorcycles.

That most of the contact events are not ‘serious’ on any industry-approved scale, does not absolve Waymo of responsibility.

On the contrary, it only offers a greater truth: Waymo vehicles, as understood through the public release of such data, are unable to cope with the everydayness of everyday environments—the normal interactions usually encountered by normal drivers in normal situations. Yet, a large part of the reason why none of these incidents are deemed serious or complex, is that as seen before, Waymo operators take control before such incidents occur, seeking to subsequently model them in software instead. Waymo vehicles have not been involved in serious incidents because Waymo operators prevent them from happening in the first place. What remains are the background bumps and scrapes of everyday interactions.

Yet, through this insistent categorization of actual contact events, Waymo is unable, or unwilling, to apportion any type of blame to their own vehicles. Not a single (minor) contact event in the 20 documented incidents is understood by Waymo to have been caused by one of their vehicles. In a world where bumps and scrapes are normal, responsibility and blame are usually offered around, divided up amongst road users and the universe at large. Yet, in Waymo’s world, Waymo is never to blame.

Deferred Decisions

To further understand the significance of this demonstration work, I want to draw both on Louise Amoore’s idea of the ‘weight of algorithms’ (Amoore, 2020, p. 163) and her concept of ‘algorithmic doubt’ (Amoore, 2020, p. 147). Regarding algorithmic ‘weight’—that algorithmic decisions must bear the weight, or burden, of what is being decided—it seems sensible to suggest that Uber ATG’s software did not or *could not* bear such weight, discounting and devaluing other road users such as Herzberg (remember she was classified in a category literally referred to as ‘Other’). Pek et al. (2020) specifically attempt to find a *fail-safe standard* so decisions made by autonomous vehicles do indeed bear the weight of the decisions they make.

Waymo—in releasing the data on contact events—is attempting to *demonstrate* how its own vehicles bear the weight of the decisions it makes in multiple ways: through use of a standardized injury severity scale, through the granular categorization of different types of contact events, through the narrative description of common conflict groups, and most significantly, through the calculation of counterfactual simulations. In this, Waymo talk of ‘putting different weights on collisions’ (Victor et al., 2023, p. 10), in order to ‘encourage the reader to consider the nature of each event and the contributions each party made to that event’ (Victor et al., 2023, p. 10). As I’ve argued above, however, this weight is never in fact borne on Waymo themselves.

On the notion of algorithmic ‘doubt’ I want to diverge a little from Amoores argument in which she suggests, in reference to machine learning, that ‘at the instant of the actualization of an output signal, the multiplicity of potentials is rendered as one, and the moment of decision is placed beyond doubt’ (Amoores, 2020, p. 137). While she goes on to discuss how ‘doubt’ is or isn’t incorporated into ML processes—she says that decisions regarding whether an entity is or isn’t a cat simply ‘embodies the truth telling of ... ground truth data of what a cat can be’ (Amoores, 2020, p. 137) rather than in any factual sense what a cat *is*—I want to return to the question of *actualization*. In the way Amoores considers it, the moment of the decision is the point at which all contingencies stop, and the ‘vast multiplicity of possible pathways’ (Amoores, 2020, p. 137) is narrowed down to, and consolidated into, a ‘single output’ (Amoores, 2020, p. 137). But as I’ve hopefully drawn out from this chapter, this narrowing of possible paths doesn’t really stop there. There is of course still a discrete, decision-made moment, but rather than drawing a close to proceedings it merely opens up another vast set of possible paths, the ‘sliding doors’ moment, or what Erin McElroy (2021, p. n.p.) refers to as the ‘flourishing’ of algorithmic trees and ‘forging’ of multiple branching points, that yields the possibility for a host of simulated contact events to be, or not be, generated—even with the use of the deterministic models I’ve just referred to.

I understand these decisions as *deferred* decisions in which the outcome of these decisions is not—and never—finalized but always already forms the possible starting point for subsequent simulations, being

pushed along a chain of calculations. As Sprenger (2021, p. 165) has suggested, in relation to autonomous vehicles: ‘a decision is never absolute – it always implies an alternative’. However, here such deferred decisions can still be distinguished from confirmatory decisions or decisions as recommendations more typical in situations where the results of an algorithmic decision are presented to a human, say, a border officer or local government employee. Whilst the making of decisions might imply openness, options, and available choices, it is clear in the case of confirmatory decisions, there is a considerable lack of choice or at least a substantial weighting to desired, ‘recommended’ decisions supported by black-boxed algorithmic calculations. What Sprenger (2021, p. 170) refers to as a ‘decisional determinism’ is variously, and unequally, at play across these instances.

In the Waymo case, these simulations are brought to bear again on the software itself. The endlessly deferred ‘non-decisive decisions’ that characterize the counterfactual simulations—perhaps, *non-decisions*—simply serve to defer their own *meaning* if not execution, per se. Not only are they simulated, that is that they didn’t happen, but they only ever serve as inputs for future decisions. In the case of the rider-only operations data, rather than being deferred per se, Waymo decides to place the decision-making of their own vehicles beyond doubt, even if decidedly doubtful. Here, doubt serves as a manageable, tangible object thrown about by Waymo in order to apportion blame. Accordingly, the Waymo vehicle is portrayed as a hapless victim, subject to the mistakes and errors of fallible decision-makers all around them.

In this, doubt is always ‘suspended’ or carried with the data, with the decision never, ultimately, being made. As Amore (2020, p. 139) writes in regard to a neural network algorithm, they are doubtful ‘not only in the sense that it supplies a contingent probability for an absolute decision, but moreover that it actively generates thresholds of normality and abnormality’. Here, in contrast, both computational process and decision are weighted and balanced, with doubt carried along both, rather than allayed. In short, it is through the simulation and classification of contact events Waymo is able to suspend, and importantly to *manage*, their (doubtful) meaning, marshalling the way in which Waymo vehicles are considered safe and careful users of the road. Arguably, this management

of meaning through the calculation of counterfactuals can be considered as an example of Enlightenment thinking, in which an unknown phenomenon is mastered, 'not so much by eliminating it but by *controlling* it' (Bates, 2002, p. 10 quoted in Aradau & Blanke, 2021, p. 4, authors' emphasis). This demonstrates, thusly, the role counterfactuals play in managing the perception of autonomous vehicles in the public realm; one of many 'problematizations of error' (Aradau & Blanke, 2021, p. 2) that now circulate in relation to machine learning and autonomous systems.

Conclusion

In Waymo's world, Waymo is never to blame. In this chapter I have considered how Waymo uses a series of techniques, built around computer simulations, to manage possible, and actual, crashes involving their own autonomous vehicles. Through an analysis of two public reports released by Waymo in 2020 and 2023, derived from autonomous vehicle data, I have argued that Waymo seeks to suspend or defer the meaning of these crashes, euphemistically referred to as 'contact events'. In the first instance, Waymo used computer simulations to model what might have happened after a human operator takes control during a test drive, where a contact event was imminent. In modelling the outcomes Waymo is ultimately able to control what might have happened, avoiding both actual harm to road users and reputational harm to the company. By establishing a narrower set of parameters in such simulations, Waymo is able to limit the range of outcomes deemed possible.

In the second instance, Waymo has developed a comprehensive set of categories to classify actual contact events involving their autonomous vehicles during so-called rider-only operations. By classifying each event according to an exhaustive typology of 'conflict groups', coupled with assigning industry-standard injury scores to each incident, Waymo seeks to apportion blame away from Waymo onto other road users and objects. In short, once again, to manage the meaning of crashes or incidents their vehicles have been involved in.

In releasing these reports, I understand Waymo as attempting to demonstrate that their autonomous vehicles are not only safe but bear the ‘weight’ of their own decision-making (Amoore, 2020). Yet, through this simulation and classification work, I argue that Waymo instead generates *deferred decisions*—non-decisive decisions that simply serve to defer their own outcomes. In so doing, they defer the meaning that such decisions are ultimately designed to hold, using algorithmic doubt as a tool with which to apportion blame away from Waymo, away from their vehicles, and onto anyone—or anything—else.

Every part of this operation—from the setting of path trajectory variables to the discussion of conflict patterns—is designed to demonstrate that Waymo is different: different from Uber ATG, different from other human drivers, and different from any contemporary autonomous vehicle operator. In establishing this difference, predicated on Waymo demonstrating their higher safety standards in conducting and evaluating their autonomous vehicle operations, Waymo asserts its superiority. Much like a casino, in Waymo’s world, Waymo always wins.

References

- Agre, P. (1994). Surveillance and capture: Two models of privacy. *The Information Society*, 10(2), 101–127. <https://doi.org/10.1080/01972243.1994.9960162>
- Amoore, L. (2020). *Cloud ethics: Algorithms and the attributes of ourselves and others*. Duke University Press. <https://doi.org/10.2307/j.ctv11g97wm>
- Aradau, C., & Blanke, T. (2021). Algorithmic surveillance and the political life of error. *Journal for the History of Knowledge*, 2(1), 1–13. <https://doi.org/10.5334/jhk.42>
- Bates, D. W. (2002). *Enlightenment aberrations: Error and revolution in France*. Cornell University Press.
- Bucher, T. (2018). *If... Then: Algorithmic power and politics*. Oxford University Press. <https://doi.org/10.1093/oso/9780190493028.001.0001>
- Culver, G. (2018). Death and the car: On (auto)mobility, violence, and injustice. *ACME: An International Journal for Critical Geographies*, 17(1), 144–170. <https://acme-journal.org/index.php/acme/article/view/1580>
- Deleuze, G. (1981). *Difference and repetition*. Columbia University Press.

- Disley, I. J. (1997). A life in OR. *Journal of the Operational Research Society*, 48(1), 3–9. <https://doi.org/10.2307/3009937>
- Elliot, G. D., & Bond, J. A. (1959). *Practical ironmaking*. United Steel.
- Ganesh, M. I. (2017). Entanglement: Machine learning and human ethics in driver-less car crashes. *A Peer-Reviewed Journal About Machine Research*, 6(1), 76–87. <https://doi.org/10.7146/aprja.v6i1.116013>
- Ganesh, M. I. (2022). *The ethical apparatus: The Material-discursive shaping of ethics, autonomy, and the driverless car*. PhD thesis. Lüneburg: Leuphana University. https://pub-data.leuphana.de/files/1348/Dissertation_2022_Ganesh_Ethical.pdf
- Gibson, J. J., & Crooks, L. E. (1938). A theoretical field-analysis of automobile-driving. *The American Journal of Psychology*, 51(3), 453–471. <https://doi.org/10.2307/1416145>
- Hawkins, A. J. (2020, October 30). Waymo pulls back the curtain on 6.1 million miles of self-driving car data in Phoenix. *The Verge*. <https://www.theverge.com/2020/10/30/21538999/waymo-self-driving-car-data-miles-crashes-phoenix-google>
- Hind, S. (2019). Digital navigation and the driving-machine: Supervision, calculation, optimization, and recognition. *Mobilities*, 14(4), 401–417. <https://doi.org/10.1080/17450101.2019.1569581>
- Hind, S. (2022). Machinic sensemaking in the streets: More-than-lidar in autonomous vehicles. In G. Rose (Ed.), *Seeing the city digitally: Reconfiguring temporalities, spatialities and collectivities* (pp. 57–80). Amsterdam University Press. <https://doi.org/10.2307/j.ctv2j6xrs3.6>
- Hollocks, B. (2006). A glimpse of early OR – A brief history of Cybor House. *OR Insight*, 19, 18–25. <https://doi.org/10.1057/ori.2006.4>
- Hong, S.-h. (2022). Predictions without futures. *History and Theory*, 61(3), 371–390. <https://doi.org/10.1111/hith.12269>
- Journal of the Operational Research Society (JORS). (2022). Scope. *Journal of the Operational Research Society*. Retrieved October 1, 2023, from <https://www.theorsociety.com/publications/journals/jors/>
- Knorr-Cetina, K. (2009). The Synthetic situation: Interactionism for a global world. *Symbolic Interaction*, 32(1), 61–87. <https://doi.org/10.1525/si.2009.32.1.61>
- Korosec, K. (2020, December 7). Uber sells self-driving unit Uber ATG in deal that will push Aurora's valuation to \$10B. *TechCrunch*. <https://techcrunch.com/2020/12/07/uber-sells-self-driving-unit-uber-atg-in-deal-that-will-push-auroras-valuation-to-10b/>

- Latour, B. (1988). *The pasteurization of France*. Harvard University Press.
- Marres, N. (2020). For a situational analytics: An interpretive methodology for the study of situations in computational settings. *Big Data & Society*, 7(2), 1–16. <https://doi.org/10.1177/2053951720949571>
- McElroy, E. (2021, February 1). Cloud ethics, review by Erin McElroy. *Society + Space*. <https://www.societyandspace.org/articles/cloud-ethics-review-by-erin-mcelroy>
- McKinlay, A., & Wilson, J. (2012). ‘All they lose is the scream’: Foucault, Ford and mass production. *Management & Organizational History*, 7(1), 45–60. <https://doi.org/10.1177/1744935911427219>
- Pek, C., Manzinger, S., Koschi, M., & Althoff, M. (2020). Using online verification to prevent autonomous vehicles from causing accidents. *Nature. Machine Intelligence*, 2, 518–528. <https://doi.org/10.1038/s42256-020-0225-y>
- Schwall, M., Daniel, T., Victor, T., Favarò, F., & Hohnhold, H. (2020). Waymo public road safety performance data. Preprint. <https://doi.org/10.48550/arXiv.2011.00038>
- Shapin, S., & Schaffer, S. (1985). *Leviathan and the air-pump: Hobbes, Boyle, and the experimental life*. Princeton University Press.
- Shields, R. (2002). *The virtual*. Routledge.
- Sprenger, F. (2020). Microdecisions and autonomy in self-driving cars: Virtual probabilities. *AI & Society*, 37, 619–623. <https://doi.org/10.1007/s00146-020-01115-7>
- Sprenger, F. (2021). Intervals of intervention: Micro-decisions and the temporal autonomy of self-driving cars. In A. Volmar & K. Stine (Eds.), *Media Infrastructures and the politics of digital time: Essays on hardwired temporalities* (pp. 157–174). Amsterdam University Press. <https://doi.org/10.2307/j.ctv1xcxr3n.12>
- Steinhoff, J., & Hind, S. (2024). Simulation and the reality gap. *MediArXiv*. <https://doi.org/10.33767/osf.io/np3vb>
- Tocher, K. D. (1960). An integrated project for the design and appraisal of mechanized decision-making control systems. *OR*, 11(1–2), 50–65. <https://doi.org/10.2307/3007180>
- Tocher, K. D., Owen, D. G., & Cunningham-Green, R. A. (1959). *Handbook for General Simulation Programme, Vol. 1*. Sheffield: Department of Operational Research and Cybernetics. Retrieved December 8, 2023, from <https://d.lib.ncsu.edu/collections/catalog/mc00653-001-lb0001-006-001#c=&m=&cv=&xywh=-261%2C3706%2C5895%2C4200>

- Uber. (2020, September 1). *Safety report 2.0: A principled approach to safety*.
Uber. Retrieved October 1, 2023, from <https://uber.app.box.com/v/UberATGSafetyReport>
- Victor, T., Kusano, K., Gode, T., Chen, R., & Schwall, M. (2023, February 28).
Safety performance of the Waymo Rider-Only Automated Driving System at
one million miles. *Waymo*. Retrieved October 1, 2023, from <https://storage.googleapis.com/waymo-uploads/files/documents/safety/Safety%20Performance%20of%20Waymo%20RO%20at%201M%20miles.pdf>
- Waymo. (2023, February 28). First million rider-only miles: How the Waymo
Driver is improving road safety. *Waymo*. <https://waymo.com/blog/2023/02/first-million-rider-only-miles-how.html>



6

Securing Decisions: Sovereignty and Semiconductors

This chapter offers a different perspective on autonomous driving. As I established in the introduction, this book is neither about the automotive industry nor the car, *per se*. To restate those arguments, car manufacturers have not, broadly considered, been interested in developing autonomous vehicles but in designing new vehicle models with assistive features, such as advanced driver-assistance systems (ADAS). Such decisions are, I argue, driven by historic, valued partnerships between automotive manufacturers and trusted suppliers. Whilst this has not prevented car manufacturers from supporting autonomous vehicle projects—Ford and Volkswagen funded Argo AI, Cruise is a subsidiary of GM, Jaguar provides vehicles for Waymo—the relationships between them have been fraught, contentious, and not all long-lasting. This chapter is therefore the exception—in more ways than one—charting various efforts to ‘secure’ the material, and territorial, future of automobility.

It is divided into two sections. The first section focuses on *software* and the sovereign capacity of autonomous vehicles to authorize the making of decisions. Engendering what Frank Pasquale (2017) calls ‘functional sovereignty’, I argue that autonomous vehicles—through their operational, *auto-nomos*, nature—constitute novel sovereign actors. The second section focuses on *hardware* and the challenge of securing the future supply

of a prize technology: the semiconductor chip. This second section argues that the re-unification of Germany in 1990 created favourable conditions for German automotive manufacturers that are intended to continue to offer unique productive advantages in the future.

In this, the decision-making capacities of autonomous vehicles are necessarily reliant upon the development of manufacturing infrastructure to produce the global volume of semiconductor chips needed to scale. Historically, as this chapter will detail, Taiwan has emerged as a critical node in this global infrastructure. The reliance on Taiwan, and specifically one company—the Taiwan Semiconductor Manufacturing Company (TSMC)—for the fabrication of semiconductor chips, was painfully exposed during the global Covid-19 pandemic. As a result of this stated ‘chip crisis’ both the US and European Union (EU) announced landmark legislation—the CHIPS and Science Act (the US), the European Chips Act (EU)—to re-localize the fabrication of cutting-edge semiconductor chips. Such efforts are set to significantly re-shape the global production and supply of chips.

In both sections I understand the securing of decisions as *exceptional*: firstly, in relation to Giorgio Agamben’s notion of the ‘state of exception’ (Agamben, 2005) and the noted exceptionality of autonomous vehicle software through which their functional sovereignty is exercised. Then secondly, in relation to the exceptional circumstances of the re-unification of Germany, the concomitant rise of TSMC and the re-integration of the former German Democratic Republic (GDR) into the Federal Republic of Germany (FRG). These two forms of exceptionality, I argue, result in the specific ‘securitization’ of decisions and decision-making.

Exceptional Software

As Gekker and Hind have written, ‘the “nomos” in autonomous suggests the vehicle is a sovereign power’ (Gekker & Hind, 2019, p. 1428). ‘Autonomous’ is derived both from the Greek for self or person (*Autos*) and from the Greek for ‘law’ or ‘custom’ (*Nomos*). Auto-*nomos* thus ordinarily refers to the power or ability to make decisions without appeal to a higher authority, like a sovereign actor such as a nation state

or monarch.¹ Auto has also become both a formal (in German, *Auto*) and colloquial (US, 'Auto') term to refer to an automobile, a term referencing the motor car's 'self-moving' ability, as opposed to being pulled by a horse. Car is likewise a truncation of 'carriage', as in 'horse and carriage'.

A nominally autonomous vehicle thus has two senses of self: firstly, in that it is self-moving (auto-mobile) and secondly, in that it is self-governing (auto-*nomos*). Of course, ordinary vehicles are described as being the former without being the latter, thus requiring a human to 'govern' it and enable its self-movability. An auto-*nomos* vehicle might also be imagined to have the capacity to exercise authority over movement but lack the actual ability to be auto-mobile (say, an autonomous carriage issuing commands to a horse pulling it). Yet, as this section unpacks, stating that a vehicle has the ability to make decisions autonomously—that is without appeal to a higher sovereign actor—is somewhat problematic, much like the notion of an auto-mobile increasingly reliant upon cloud infrastructures and wireless connectivity to enable self-movement.

As Benjamin Bratton writes:

The implication is not that software is new and sovereignty is timeless, thereby leading one to ask how sovereignty now works through software, but rather that both are now mutually contingent and that the work of software at a global scale itself produces *unfamiliar sorts of sovereignties*. (Bratton, 2016, p. 20, emphasis added)

The definition of auto-*nomos* as outlined previously, as 'the power or ability to make decisions without appeal to a higher authority', is thus also 'the power to suspend the regularity of the law' (Bratton, 2016, p. 20), following Giorgio Agamben's (Agamben, 2005) reading of the conservative German jurist Carl Schmitt. Here, *nomos* is understood as 'a making of a territorial order through the execution of a territorial claim and physical occupation that precedes it' (Bratton, 2016, p. 25). *Nomos* is seen by Schmitt as demonstrating a particular geopolitical dimension,

¹Throughout this chapter I sometimes use the formulation 'auto-*nomos*' to foreground these etymologies.

depending on who or what is executing a territorial claim. The European *nomos*, thus, is seen differently from a US *nomos* in which territorial claims and physical occupation of space occur, and have occurred, historically differently. As Bratton contends, ‘today the continuing (if still incipient) emergence of planetary-scale computation may represent a similar break and a similar challenge to the political geographic order’ (Bratton, 2016, p. 26) as the partition of the New World or similar colonial periods of history (i.e. the ‘Scramble for Africa’).

For Schmitt (1985 [1922]), a sovereign actor is ‘he who decides on the state of exception’ (quoted in Agamben, 2005, p. 1). As Agamben suggests, such a sovereign actor is not somehow outside of or beyond the law, ‘because the state of exception represents the inclusion and capture of a space that is neither outside or inside’ (Agamben, 2005, p. 35). Quoting from Schmitt (1985 [1922]), Agamben writes that the ‘sovereign stands outside [*steht außerhalb*] of the normally valid juridical order and yet belongs [*gehört*] to it, for it is he who is responsible for deciding whether the constitution can be suspended *in toto*’ (Agamben, 2005, p. 35). Thus, the sovereign actor decides whether normal order applies, or whether there are special conditions in which normal order has been expended, and exceptional rules apply.

This has implications, naturally, for how one considers decision-making. Quoting Schmitt again, Agamben writes that ‘in suspending the norm, the state of exception “reveals [*offenbart*], in absolute purity, a specifically juridical formal element: the decision”’ (Agamben, 2005, p. 34). These decisions, thus, ‘show their autonomy’ (Agamben, 2005, p. 34) by having independent force in law. Quoting Schmitt that:

Just as in the normal situation the autonomous moment of decision is reduced to a minimum, so in the exceptional situation the norm is annulled [*vernichtet*]. And yet even the exceptional situation remains accessible to juridical knowledge, because both elements, the norm as well as the decision, remain within the framework of the juridical [*im Rahmen des Juristischen*]. (Agamben, 2005, p. 34)

The question is the extent to which the autonomous vehicle might be understood as exercising a continual possible state of exception that, as a

sovereign actor, it is capable of enacting at any given moment. What is crucial here is understanding the extent to which internally logical decisions are executed in light of obvious, or apparent, normal legal transgressions. For instance: if an autonomous vehicle were to repeatedly drive over a designated speed limit in order to avoid other potential risks calculated by the vehicle itself (but perhaps heretofore not calculated by, or enacted by, human drivers). Here, such legal transgressions are made justifiable based on the calculative efforts of the vehicle itself as a functional, machinic entity. Chapter 8, on the travails of Cruise in California, provides a concrete example of how such legal transgressions are made justifiable in the eye of the vehicle and vehicle operator.

Thus, if an autonomous vehicle generates what Bratton (2016) understands as an ‘unfamiliar sort of sovereignty’, then what does this sovereignty look like, and how does it act? If, through Agamben’s reading of Schmitt, the (familiar, ordinary) sovereign actor is able to suspend normal legal states, acting neither wholly outside of nor inside the law, then how might an autonomous vehicle act differently?

One such approach to this question would be by understanding the autonomous vehicle as utilizing a different or ‘unfamiliar’ sense of the state of exception—a question of how autonomous vehicles make sense of the world. That, rather than making any kind of final, albeit temporary decision to decree a state of exception or not (in relation to a general theatre of action—e.g. an invasion or war, or torture), that the autonomous vehicle instead sets ongoing, operational thresholds at which exceptionality is reached or not. The difference here, perhaps, is that whilst exceptionality is determined in relation to the making of decisions (or not), the ‘state’ determinable is not a geographical nor territorial one within which juridical order applies, but an operational one pertaining to specific calculative actions or trajectories. This differs, ever so slightly, from Bratton’s understanding of ‘platform sovereignty’ as enabling the drawing of reversible ‘geographic lines of segmentation ... [able to] divide physical space or separate layers in a larger machine’ (Bratton, 2016, p. 21). In an understanding of autonomous vehicles as fully sovereign actors, the lines of segmentation are not necessarily geo-territorial nor layered but protocological lines (Galloway, 2004) along which decisions flow, but also depart or diverge from one another, based on

calculative possibilities. Rather than a ‘geographic line’ (Bratton, 2016, p. 22), it is a navigational line, as the previous chapter would suppose.

For example, continuing with the speed limit case above, that an autonomous vehicle might establish thresholds of exceptionality for each calculative trajectory it generates in relation to a particular set of calculable risks. If such a vehicle decides that speeding away from a potential crash, resulting in it exceeding the speed limit, is preferable to staying within the limit but increasing the possibility of being involved in a crash, then the vehicle might be said to be operating with a sense of ‘operational exceptionality’. Whilst this form of exceptionality clearly has a spatial dimension, it is perhaps less significant that such a decision takes place in a certain space or not, and more significant that this exceptionality can essentially be ‘switched on’ or ‘switched off’ at any given moment, as thresholds are exceeded temporarily before the vehicle returns to a normal operating state, fully within the rules of the road.

The question of whether an autonomous vehicle is necessarily a truly sovereign actor—albeit in an ‘unfamiliar’ form, as Bratton suggests—is thus something only provable in relation to other possible sovereign actors, operating along more ordinarily sovereign lines. That is, whether an autonomous vehicle is ordinarily subject to the regulatory decision-making of other sovereign actors such as law enforcement. It is this aspect I turn to next.

Functional Requests

Law scholar Frank Pasquale writes of how ‘major digital firms’ are now ‘market makers, able to exert regulatory control over the terms on which others can sell goods and services’ (Pasquale, 2017, p. n.p.). Rather than simply being participants in a marketplace ordinarily governed by national or supranational regulators, they increasingly ‘aspire to displace more government roles over time, replacing the logic of territorial sovereignty with functional sovereignty’ in which citizens are ‘subject to corporate, rather than democratic, control’ (Pasquale, 2017, p. n.p.). In essence, as others have considered, that ‘big tech’ firms like Google (Alphabet) and Amazon are so powerful and have such power over the

markets they operate in, that they exert a kind of monopoly control historically akin to state actors. Leigh Phillips and Michal Rozworski extend this argument to argue that such firms likewise organize internal activities (i.e. the tracking and distribution of goods) through forms of central planning, akin to socialist state economies (Phillips & Rozworski, 2019).

Yet, following Pasquale (2017), whilst big tech firms might appear to act *like* state actors, the way they exert their power is notably different. Rather than utilizing their *de jure* (or sometimes *de facto*) right to solely govern territory and populations based, ordinarily, within legally ratified demarcated spaces, such firms exert their sovereign power through their functioning as technological mediators, governing relations between workers and work, citizens and activities. As Pasquale continues:

For example: Who needs city housing regulators when AirBnB can use data-driven methods to effectively regulate room-letting, then house-letting, and eventually urban planning generally? Why not let Amazon have its own jurisdiction or charter city, or establish special juridical procedures for Foxconn? Some vanguardists of functional sovereignty believe online rating systems could replace [US] state occupational licensure – so rather than having government boards credential workers, a platform like LinkedIn could collect star ratings on them. (Pasquale, 2017, p. n.p.)

In each of these various cases, the firms mentioned (AirBnB, Amazon, Foxconn, LinkedIn) are touted as privatized deliverers of public services, whether through the municipal regulation of a city housing market or the verification of the right to work. However, in each case, this process isn't simply a narration of the privatization of public services, but through the technological interests of each actor, that each service is provided *functionally differently*. In this, AirBnB is able to effectively regulate a city housing market by setting the conditions and costs of listing/letting properties on its platform, or LinkedIn is able to officiate labour legality through its own ratings system. This sovereign power of big tech firms, then, is not ordinarily exerted by them through control over geographical territories but through their total control over technological systems designed, managed, and operated by them.

In the case of autonomous vehicles, this functional sovereignty takes a similar form. In drawing attention to the privacy implications of allowing firms to singularly control, and manage access to, critical infrastructure, Gekker and Hind suggest that autonomous vehicles will likely ‘become novel mobile, functional, sovereign objects, through which all requests must flow, while government departments are relegated to secondary, yet likely “approved”, partners’ (Gekker & Hind, 2019, p. 1429). In this arrangement, the autonomous vehicle becomes the ‘default’ sovereign actor in any resultant claim or contestation, to which a government department or body (say, a highways agency or police force) must comply with, rather than vice versa. In many ways, this is already the case, with established big tech firms already acting as regulators for activity on their own platforms, by third parties (Van Loo, 2020; Hind & Seitz, 2024).

What is interesting in the case of autonomous vehicles, of course, is that the sovereign actor in question is not a platform, *per se*, but a mobile object ordinarily being powered by one. Whilst AirBnB’s functional sovereignty is exercised through its listings platform and various associated mechanisms enforced when users list and request properties, and Amazon’s functional sovereignty is arguably exercised through sales analytics (as actor in, and operator of, Amazon ‘Marketplace’) or computational infrastructure (as owner of cloud services provider, AWS), an autonomous vehicle manufacturers’ functional sovereignty is exercised through their control over ML-driven decision-making. Here, public authorities do not have control over what decisions human drivers make, only remit over the adjudication, and enforcement, of the legality of such decision-making. If a human driver is caught speeding by a speed camera operated by a local authority, they are usually expected to have to justify (or appeal) their doing so. Whilst there are evidently ways of resisting or countering such administrative processes (extenuating circumstances, wrong car, different driver, etc.), the role of the administrator (the state, local authority) itself is not in question, nor is the process through which a decision itself is reached.

In the case of determining the identity and reason for an autonomous vehicle committing a driving infraction is, arguably, somewhat different. In such a case, following the argument outlined above, this would involve an administrator of the legislative process likely having to submit a

request to the vehicle operator in order to verify details of the infraction. Such a situation inverts the relationship between state actor and firm, turning the question of territorial sovereignty into one of functional sovereignty. Potentially incriminating data and/or evidence collected by the vehicle, therefore, would be of great significance, in the way mobile phone data and/or app use often is already to law enforcement and insurers (Hill, 2024).² Accordingly, 'the condition of this changing relationship is that non-opt-out-able data captured through ... infrastructure are enclosed by default, with access granted by agreement or discretion, and offered as a subsequent, selective service', with political bodies assuming a new role as 'mere service users' (Gekker & Hind, 2019, p. 1429). The reason for the suspension of Cruise's operations in California lay precisely in relation to the firm failing to provide relevant video footage of an incident involving one such vehicle (Korosec, 2023), as Chap. 8 details.

The 'functional sovereignty' (Pasquale, 2017) of the autonomous vehicle is itself strengthened by a form of digital or technological sovereignty sought by nation states. In Brattonian terms, there is a soldering of the 'user layer' and the 'cloud layer' (Bratton, 2016), as the auto-*nomos* vehicle and the chip firms that nation states seek to entice, exercise their newly found sovereign power. It is this that I turn to next, as contemporary 'polycrisis' strengthens the realizable autonomy of both.

Familiar Sovereignty and Emerging Crises

The automobile as auto-*nomos* is also increasingly dependent on more familiar forms of territorial sovereignty and power. The EU's desire to secure their own digital/technological sovereignty from both the US and China, as this section will demonstrate, is representative of the growing significance of digital technologies to economic development, as well as rising global political and economic tension and instability.

²Elaine Herzberg was subject to various access requests by the National Transportation Safety Board (NTSB) for streaming services she subscribed to, including Hulu, with the NTSB naturally reliant upon both the data collecting activities and compliance of the firms in question.

One way to understand this growing tension is through the notion of ‘polycrisis’, as historian Adam Tooze has variously done in relation to the Covid-19 pandemic (Tooze, 2020), the Russian invasion of Ukraine (Tooze, 2022a), and the cost-of-living crisis (Tooze, 2022b). Polycrisis, as Tooze (2020) mentions, was a term often used by the then-European Commission President Jean-Claude Juncker. In a speech at the Annual General Meeting of the Hellenic Federation of Enterprises (SEV) in Athens, Greece, in 2016, Juncker said that he had ‘often used the Greek word “polycrisis” to describe the current situation’ (Juncker, 2016, p. n.p.). At the time, the ‘various challenges’ Juncker highlighted ranged ‘from the security threats in our neighbourhood and at home, to the refugee crisis, and to the UK referendum’ (Juncker, 2016, p. n.p.). These various challenges were not just simultaneous, Juncker said, but also ‘feed into each other, creating a sense of doubt and uncertainty in the minds of our people’ (Juncker, 2016, p. n.p.). Whilst Greece had a special place in this ongoing polycrisis, emerging as it did from the 2007 financial crisis, the EU itself (with Greece as a member state) was having to deal with this generalized, interconnected, ongoing sense of insecurity.

It is arguable whether the notion of polycrisis is at all new or special. Through the work of Antonio Gramsci, Milan Babic understands the crisis of the ‘international liberal order’ (ILO) since 2008 as occurring at three levels (global political-economic, state, societal), concerning the ‘processuality’, ‘organicity’, and ‘morbidity’ of the crisis (Babic, 2020). Bruno Amable and co-authors have likewise considered overlapping political and systemic crises in France since the 1980s (Amable et al., 2012), whilst Cédric Durand contextualizes Adam Tooze’s (Tooze, 2018) study of the aftershocks of the 2007 financial crisis throughout the subsequent decade ‘undergoing a process of “mutation and metastasis”, involving new political and geopolitical depths’ to this day (Durand, 2019, p. n.p.). In each case, there is varied weight on the interlocking nature of crises as they manifest in different economic and societal arenas, both intrinsic and extrinsic to their operation. Moreover, that the Western capitalist model has been in some kind of continued, contorting, crisis (or crises) since 2007, failing to deal with its root causes for the past 15 years.

One polycrisis response has emerged out of the confluence of the Covid-19 pandemic, the Russian invasion of Ukraine and global inflationary pressure: deglobalization. Here, deglobalization has been touted as a reaction to various ‘exogenic shocks’ manifesting from outside knowable political and economic spheres, much like a global pandemic. At various moments between 2020 and 2021, in response to the spread of Covid-19, most if not all nation states enacted some form of international travel ban or restriction, limiting arrivals into countries. Coupled with production shutdowns manifested by safe distancing protocols and localized lockdowns (especially in Shanghai, China) (Reuters, 2021), the unfurling of an interconnected political, social, and economic crisis atop of a manifest health crisis became evidently more possible. Adam Tooze (2022b, p. n.p.) counts seven ‘macroscopic risks’ including the ongoing mutation of Covid-19 variants, and the risk of economic stagflation, nuclear escalation, and the climate crisis.

Reacting to the pandemic and specific aspects of the resulting crisis I will consider later, the EU began to formulate a very particular ‘techno-protectionist’ policy direction. This is different from a broadly historical ‘techno-liberalist’ approach in the US and ‘techno-authoritarianism’ seen in China—each a distinct ‘techno-nationalist’ strategy (Rikap & Lundvall, 2021, p. 2) in itself—such that any response to polycrisis has a very distinct effect on those developing, producing, supplying, assembling, selling, and using such technology.

The background to this developmental path which, I argue, determines the extent to which the future of autonomous driving is ‘secured’ or not, is that polycrisis is fracturing the European political system to the extent that it is creating the conditions for a more interventionist EU (Lavery, 2024). As political scientists Zeitlin et al. (2019) explore, contemporary polycrisis is resulting in a so-called polycleavage leaving the EU in what they call a ‘politics trap’ (Zeitlin et al., 2019, p. 967), in which member states’ own political manoeuvring prevents higher-level (i.e. supranational/EU-level) policy-making, sometimes explicitly using EU policy to ‘mobilize’ national audiences (Zeitlin et al., 2019, p. 968) in opposition to the EU itself (Brexit, etc.).

Polycrisis as a particular manifestation of crises offers opportunities to political and economic leaders to act in ‘exceptional’ ways. Within the

EU such crises present opportunities both for greater integration between member states and greater fragmentation. In recent times, as Zeitlin et al. write:

The old functionalist adage that ‘integration advances through crises’ appears to be simultaneously confirmed and rejected: while institutional integration points in the direction predicted by neofunctionalists, the dynamics of political fragmentation have accelerated, as postfunctionalists would expect. (Zeitlin et al., 2019, p. 964)

The Russian invasion of Ukraine in 2022, for instance, has resulted in a closer relationship between NATO member states, including the rapid accession of Sweden and Finland into the military alliance itself (NATO, 2022). Yet, in contrast, with the resultant aftershocks to the global energy market, some EU member states such as Germany have been faced with greater political fragmentation and disagreement, at a regional and national levels. This fallout has significantly affected the ruling Social Democratic Party (SPD) who came to power in late 2021, despite claims to a *Zeitenwende* or historic ‘turning point’ in German fiscal and military policy, following approval of €100bn defence funding (Financial Times, 2023).

In this, polycrisis is not inherently destabilizing, with polycleavages arguably resulting in greater ‘social stability, since they distribute political divisions and grievances over a larger number of actors and policies’ (Zeitlin et al., 2019, p. 966), potentially diffusing political opposition and limiting fragmentation. Maintaining the kind of poly-centric centrism witnessed with the EU, therefore, might be compared favourably to, for instance, the ‘inherent instability of the Weimar Republic or the Austrian First Republic’ (Zeitlin et al., 2019, p. 966) characterized by marked political poles both on the left and the right. Polycrisis thus might be seen merely as an ordinary, rather than extraordinary, working environment, with equal tension between fracture and integration cancelling each other out. Germany perhaps is the most concrete contemporary example of such that despite clear fissures, its federalized system is designed precisely to dilute and neuter internal discord.

However, both the Covid-19 pandemic and the Russian invasion of Ukraine have precipitated *immediate* and *significant* economic crises deeply affecting both previously secure Northern core states, in the way neither the Eurozone crisis (Northern states as creditors, Southern as debtors) nor refugee crisis did (arguably political rather than economic), as a result of key member states' reliance upon non-EU economies. Germany, for instance, has been almost uniquely reliant upon Russian gas since its decision to close nuclear power plants in 2011 (Associated Press, 2011) and phase out coal-burning power stations by 2030 (Wacket, 2021). In recent years, Germany has also become hugely reliant—much like the global economy—on exporting goods to China. The People's Republic of China has been Germany's biggest trading partner since 2016 (Destatis, 2022a), exporting over €103billion worth of goods in 2021 (Destatis, 2022b). Automobiles and automotive parts represent 15.3% of Germany's export total in 2021 (Destatis, 2022c), much of which is exported to China. Shutdowns and production restrictions both in Germany and China during the Covid-19 pandemic thus created the conditions for the ongoing polycrisis, coupled with rising energy insecurity.

As Zeitlin et al. write, the 'specific rule set and consensus-based political system of the EU' (Zeitlin et al., 2019, p. 966) mean that crises and polycrisis can stymie political responses, unravelling instability across the bloc. Indeed, that the distinct macro-regions within the EU that Zeitlin et al. define as 'North-Western, South-Eastern and Central-Eastern' are 'in turn ... affected very differently by different components of the polycrisis' (Zeitlin et al., 2019, p. 968). Germany, with its famous manufacturing *Mittelstand* central to its economy, has also suffered more acutely than the likes of Sweden or the Netherlands, fellow North-Western member states. Central-Eastern member states such as the Czech Republic and Slovakia—significantly integrated into Germany's manufacturing supply chains as junior, supply side, partners, have suffered because of the arrangement. As journalist Marco D'Eramo has written, 'Germany has sought to construct a series of mutually interdependent economies which now essentially amount to a single economic system' (D'Eramo, 2022, p. n.p.) with the aforementioned Eastern member states of the Czech

Republic and Slovakia ‘seats of the automobile industry’ (D’Eramo, 2022, p. n.p.).

In response to the effects of a globally derived polycrisis, therefore, regional blocs such as the EU as well the US are seeking to develop increasingly independent economic ‘stacks’, able to ride out or wholly prevent various crises, and avoid supply chain, security, and material extraction issues.

In this next section I examine how the growing use of semiconductor chips in the automotive industry has created the conditions for one such crisis, the 2022 chip crisis.

Automotive ‘Chipification’

The semiconductor chip crisis significantly affected the automotive industry because of the 50-year ‘chipification’ of automobiles (Forelle, 2022). As Forelle writes, ‘chipification is radically reshaping such processes as resource allocation, labor flows, and cultural practices around car manufacture, use, repair, and modification’ (Forelle, 2022, p. 1). Whilst automobiles have contained semiconductor chips since the 1960s, the platformization of cars (Hind & Gekker, 2022; Hind et al., 2022) has led to a significant increase in the number of chips per vehicle. As a *New York Times* article suggests, a ‘modern car can easily have more than 3,000 chips’ (Ewing & Boudette, 2021, p. n.p.), controlling everything from fuel intake to electric windows.

The Taiwan Semiconductor Manufacturing Company (TSMC), the world’s leading chip fabricator, considers that the acceleration of the ‘adoption of electric vehicles (EVs), advanced driver-assistance systems (ADAS) and smart cockpit/infotainment systems, along with new electrical/electric (E/E) architecture’ (TSMC, 2022, p. 18) will drive ‘increased demand for AP/MCU/ASIC processors, in-car networking, sensors, and power management ICs [integrated circuits], thus continuously increasing the silicon content per car’ (TSMC, 2022, p. 18).

Forelle (2022) understands chipification from three perspectives: global supply chains, labour, and car cultures. In the first instance, the rise of semiconductor chips in automobiles has put greater focus on the

global production and sourcing of not only finished chips but raw materials integral to the production of the chips themselves. Semiconductor chips have historically been dependent upon the sourcing of silicon and germanium, ‘obvious examples of discoveries in chemistry that proved to be essential for computer culture’ (Parikka, 2015, p. 36), as well as, increasingly, car culture. The production of the alloy silicon germanium in the 1980s by researchers at IBM led to significant increases in chip efficiency and performance, as well as cost savings, compared to silicon-based chips alone (IBM, 2022). As Parikka further writes:

Transistor-based information technology would not be thinkable without the various meticulous insights into the material characteristics and differences between germanium and silicon – or the energetic regimes – whether that involves the consideration of current clouds (as in server farms) or the attempts to manage power consumption inside computer architectures. (Parikka, 2015, p. 57)

Sourcing raw material for, and manufacturing, semiconductor chips requires huge amounts of energy. Whilst silicon itself is abundant on earth, it still requires an extensive process to extract it from where it is found, like quartz rock. As journalist Douglas Heaven has explained:

Rocks extracted from the ground with machines and explosives are put into a crusher, which spits out quartz gravel. This then goes to a processing plant, where the quartz is ground down to a fine sand. Water and chemicals are added to separate the silicon from other minerals. The silicon goes through a final milling before being bagged up and sent as a powder to a refinery. (Heaven, 2019, p. n.p.)

Once the silicon is reduced to a powder, it is sent to chip fabrication facilities commonly referred to as ‘fabs’ or ‘foundries’ where:

The [silicon] material is melted in a furnace at 1,400°C and formed into cylindrical ingots. These are then sliced into discs called wafers, like chopping up a cucumber. Finally, several dozen rectangular circuits – the chips themselves – are printed onto each wafer in factories ... From here, chips make their way to every corner of the planet. (Heaven, 2019, p. n.p.)

Only once quartz rock has been mined, silicon has been separated from other minerals, then milled, melted, reformed, sliced, printed, and shipped, do semiconductor chips finally make their way to the car manufacturer assembly lines, usually via key, ‘tier one’, suppliers like Bosch tasked with integrating them into vehicle modules like infotainment systems. For this to happen, car manufacturers have historically been reliant upon global semiconductor chip supply chains.

As Bratton writes, ‘there is no planetary-scale computation without a planet, and no computational infrastructure without the transformation of matter into energy and energy into information’ (Bratton, 2016, p. 75). Chipification, thus, has taken the automotive industry down a path that intensifies planetary-scale computation, increases automotive manufacturers’ reliance upon semiconductor chips, and multiplies manufacturers’ exposure to geopolitical shocks and crises. What is perhaps most intriguing about such a situation is that car manufacturers are increasingly pressured to execute a double movement precipitated by chipification.

Hard Software, Soft Hardware

On the one hand, manufacturers are seeking to develop slicker, integrated, ‘platformized’ software systems to ‘datafy’ the driving experience (Hind, 2021), driven by wider innovations in the tech industry (Hind et al., 2022). German automotive manufacturer Volkswagen, especially keen to innovate in this field, have suffered considerable issues, with owners of vehicles from their ID electric models, reporting ‘problems with infotainment screens, range calculations, buggy smartphone connections, charging, and other features that are far more seamless on other companies’ cars’ (George, 2022, p. n.p.). Whilst these problems manifest at the interface for the driver, they begin with the chips themselves, through which, as a Volkswagen spokesperson suggests, ‘software will become the new differentiator’ (George, 2022, p. n.p.). For car companies, this means crossing into new semiconductor chip markets able to secure the kinds of chips required to power infotainment systems and next-generation

dashboards, markets dominated by mobile device manufacturers like Apple, Samsung, and Huawei.

For Volkswagen's software arm, Cariad, this also requires the 'future-proofing' of hardware, capable of hosting an 'exceptionally powerful and scalable computing platform' over the long term (Cariad, 2023, p. n.p.). One such aspect of Cariad's stated approach is the 'decoupling of software and hardware' (Cariad, 2023, p. n.p.), such that Volkswagen can account for the different innovation and iteration cycles between each. Another aspect focuses on the components themselves, grouping and limiting their overall number to reduce complexity. Whilst Cariad is referring to vehicle operations, it also patently applies to Volkswagen's own supply chains: reducing component complexity means fewer chips and a more streamlined supply chain. Yet, with glitchy innovations at Volkswagen being rolled back, such as capacitive steering wheel controls being switched for physical buttons (Gitlin, 2022), it is clear that car manufacturers are struggling with both software and hardware at all scales (Carr & Welch, 2024).

Chip Production

TSMC is one of the world's largest chip manufacturers, responsible for supplying semiconductor chips to the likes of AMD, Apple, and Nvidia. Indeed, Taiwan is a leader in the production of semiconductors, followed by South Korea, Singapore, and China. The EU, in comparison, possesses less than a 10% share of global chip manufacturing facilities.

TSMC was founded in 1987 by Morris Chang, a former senior executive at Texas Instruments (TI), a Dallas, Texas-based technology firm involved in the fabrication of semiconductor chips. Whilst born in Ningbo, China, Chang received his PhD from Stanford before moving to Texas Instruments. By the mid-1980s Taiwan—long a critical location in the *assembly* of semiconductor chips—realized it needed to develop its own fabrication base, in part, because of China's entry into the global chip assembly market (Miller, 2022). Moving up the supply chain was seen as the only way to avoid direct competition with the emerging superpower. In 1985, Taiwan's economic minister Kwoh-ting Li convinced

Chang to set up TSMC, with the Taiwanese government providing 48% of the initial capital, 'stipulating only that Chang find a foreign chip firm to provide advanced production technology' (Miller, 2022, p. 167). Despite TI and Intel rejecting their advances, Dutch semiconductor firm Philips committed \$58million to 'transfer its production technology, and license intellectual property in exchange for a 27.5% stake in TSMC' (Miller, 2022, p. 167).

Since 1985 TSMC has 'grown into a giant with an effective stranglehold on the global chip supply chain' (Hille & Sevastopulo, 2022, p. n.p.) such that it effectively serves as a 'critical security guarantee' (Hille & Sevastopulo, 2022, p. n.p.) in the face of possible aggression from China, constituting a so-called silicon shield (Hille & Sevastopulo, 2022, p. n.p.) for the country. Indeed, crucial to understand are the 'deep ties' (Miller, 2022, p. 167) TSMC specifically has with the US chip industry, long constituting a 'symbiosis' that has 'benefitted Taiwan and Silicon Valley' (Miller, 2022, p. 168). Many of TSMC's hires came from the US and even from TI specifically, whilst 'throughout much of the 1990s, half of TSMC's sales were to American companies' (Miller, 2022, p. 168). In 2022, 68% of TSMC's \$75.88billion consolidated net revenue came from US-based customers (TSMC, 2022, p. 18). According to their own figures TSMC holds a 30% global market share in the 'foundry' (i.e. chip fabrication facility) segment of the semiconductor chip industry in 2022, up from 26% in 2021 (TSMC, 2022).

Many leading semiconductor companies like Qualcomm, Nvidia, and AMD are 'fabless', outsourcing fabrication to third parties, such as TSMC. This means that even fewer companies are responsible for producing the semiconductor chips required for everything from washing machines to video games consoles. Volkswagen, for example, rely on the Bavarian semiconductor manufacturer Infineon, to supply them with chips. Their new electric ID.4 model, central to their new electric vehicle strategy following the diesel emissions scandal in 2015, contains 50 such Infineon chips, from power semiconductors to microcontrollers (Fine, 2021). Other European chip companies include Bosch (also Germany), STMicroelectronics (France/Italy), and NXP (Netherlands). The automotive sector is considered by the likes of TSMC and others as a distinct chip market segment itself, alongside smartphones, high-performance

computing (PCs, tablets, games consoles, etc.), Internet of Things, and ‘digital consumer electronics’ (TVs, set-top boxes, etc.) (TSMC, 2022, p. n.p.).

Whilst contemporary vehicles require different types of semiconductor chips, none need a ‘wafer capacity’ (Kleinhans, 2021, p. 9) of 10 nanometres (nm) or below—the current ‘cutting-edge’ node size of semiconductor chips. Chips at 10 nm and now 7 nm are ordinarily being fabricated for flagship mobile devices, such as Apple iPhones or iPads. Apple’s iPhone XS A12 Bionic system on a chip (SoC), launched in 2018, was the first such mass market device to use 7 nm nodes, both faster (by 15%) and more power efficient (by 40%) than a previous edition of the iPhone X (2017) (Shankland, 2018). What this means, of course, is that it is major mobile device manufacturers rather than car manufacturers who are driving advanced wafer fabrication. As the German thinkthank *Stiftung Neue Verantwortung* (SNV)³ write:

Node shrinkage or ‘More Moore Scaling’ (more transistors per square millimeter with better performance, less power consumption and lower costs) is especially important for logic semiconductors in the consumer market, such as processors in smartphones and laptops. At the same time, this market segment generates the high volumes that make the enormous upfront investments in chip design and manufacturing economically viable: For example, Apple’s iPhone sales totaled US\$65.6 billion in 4Q 2020 alone. With that amount of quarterly sales, Apple can afford to invest billions in its own chip design and contract with TSMC to manufacture the chips. Because of these economies of scale, the most advanced chips in terms of ‘PPAC’ (power, performance, area, cost) are found in consumer electronics. (Kleinhans, 2021, p. 8)

The production of semiconductor chips is complex, hugely expensive, and globally distributed, like many electronics industries (Yeung, 2022). US tech firms like Apple rely on semiconductors manufactured in Taiwan, but TSMC do not design the chips themselves, Apple does. As Chris Miller (2022, p. xxiv) illustrates:

³ In English, New Responsibility Foundation.

A typical chip might be designed with blueprints from the Japanese-owned, UK-based company called Arm, by a team of engineers in California and Israel, using design software from the United States. When a design is complete, it's sent to a facility in Taiwan, which buys ultra-pure silicon wafers and specialized gases from Japan. The design is carved into silicon using some of the world's most precise machinery...produced primarily by five companies, one Dutch, one Japanese, and three Californian...Then the chip is packaged and tested, often in Southeast Asia, before being sent to China for assembly into a phone or computer.

Yet, this wasn't always the case. The rise of fabless chip firms is concomitant with the rise of TSMC:

Before TSMC, a couple of small companies, mostly based in Silicon Valley, had tried building businesses around chip design, avoiding the cost of building their own fabs by outsourcing the manufacturing...Not having to build fabs dramatically reduced startup costs, but counting on competitors to manufacture chips was always a risky business model. (Miller, 2022, p. 168)

Chris Miller likens this development to the invention of the Gutenberg Press in fifteenth-century Germany, with fabless chip firms embodying a 'democratization of authorship' (Miller, 2022, p. 168) in the broad ability for such firms to design their own chips with underlying technology. Yet, coupled to this was the concentration, consolidation, and centralization of chip fabrication itself, thanks to the capital-intensive nature of the manufacturing process. Yeung (2022) characterizes this as a three-fold development, combining a 'geographical shift in wafer fabrication toward far greater concentration in East Asia' (Yeung, 2022, p. 129) with an 'industrial-organizational shift from IDM [Integrated Device Manufacturer] firms worldwide to foundry service providers mostly based in East Asia' (Yeung, 2022, p. 129) and an 'end-market shift ... toward rapidly growing ICT segments' (Yeung, 2022, p. 129) such as the automotive industry.

Together with the famously complex supply chain of automotive manufacturers, and semiconductor chip production becomes one fragile, susceptible manufacturing process bolted onto another. 'Unlike oil' as Miller

continues, ‘production of computing power depends fundamentally on a series of choke points: tools, chemicals, and software that often are produced by a handful of companies – and sometimes only by one’ (Miller, 2022, p. xxv) such as TSMC, or ASML, the Dutch chip machine firm referenced above by Miller, or Japan’s SCREEN, relied on for chemical cleaning systems (Ting-Fang & Li, 2022).

The situation is similar higher up the supply chain too, where the extraction of raw materials is no simpler and cost-effective than building complex machinery. As Ting-Fang and Li write, ‘follow the supply chain upstream, and further chokepoints emerge with regard to the fluoropolymers from which [chip] components are made’ (Ting-Fang & Li, 2022, p. n.p.). Only a small number of fluoropolymer material producers exist, and some materials like PFA ‘require extensive knowhow to process’ with ‘no competitors on the horizon’ (Ting-Fang & Li, 2022, p. n.p.) besides a US firm (Chemours) and a Japanese company (Daikin Industries) to make them.

Higher still, and fluoropolymers are processed from a mineral called fluorspar or fluorite, ‘often labelled as a “semi-rare earth”’ mineral (Ting-Fang & Li, 2022, p. n.p.) whose production is dominated by China (60%). Whilst other producers like Mexico (10.8%), Mongolia (8.2%), and South Africa (4.5%) can also be relied on, the semiconductor industry is one amongst many that rely on the processing of fluorspar, eventually used as a resin or coating for the valves, pumps, tubes, pipes, and containers integral to chipmaking equipment and cleaning systems (Ting-Fang & Li, 2022). In other words, it is easier to list the components that are readily available than those that aren’t, in the chip manufacturing process.

Under ordinary operating conditions the fragile semiconductor supply chain just about holds together. Under extraordinary conditions everything, as the last few years have shown, falls apart. This has resulted in a suite of historically unparalleled efforts by major nations (the US, China) and supranational blocs (EU) to shore up the semiconductor supply chain, and to prevent future blockages in the manufacture of chip-reliant consumer electronics. I document these efforts to ‘engineer’ the chip crisis, with a specific focus on the EU, in the next section.

(Engineering the) Chip Crisis

Without high-quality semiconductor chips, the computational capacity required for vehicles to process image and sensor data and execute ML-driven processes is impossible. With the Covid-19 pandemic, this vulnerability intensified, with supply chain issues from key suppliers resulting in a global shortage of chips, and a slowdown in the production of vehicles. A resultant scramble to increase fabrication capacity in the US and the EU, however, is perhaps unlikely to offer any solutions. SNV have suggested that the historic lack of investment in ‘cutting-edge’ semiconductor fabs has left the EU with weakened demand (Kleinhans, 2021).

Connectedly, they ask why fabless US firms (i.e. Qualcomm, AMD) would choose production facilities based in the EU over established foundries in Taiwan or South Korea. Lastly, they argue that ‘skyrocketing’ (Kleinhans, 2021, p. 9) investment costs (US\$20 billion/fab), R&D ‘intensiveness’ (Kleinhans, 2021, p. 9), and the need for near total production capacity utilization (90%+) have led to many firms exiting the semiconductor market altogether. Only ‘Samsung in South Korea and TSMC in Taiwan’, they suggest, are equipped to deliver ‘cutting-edge fabs’ (Kleinhans, 2021, p. 2).

Two relevant observations can be made regarding the above. Firstly, that car manufacturers do not possess the capital to make similar such investments, although, naturally, big tech firms do. This is in part because most car manufacturers are currently investing their precious capital in the electrification of drivetrains, shifting away from internal combustion engines. Within the EU, this effort has been driven by legislation mandating a ban on petrol and diesel vehicles from 2035 (European Parliament, 2022). Such a relatively hard deadline has required manufacturers to dedicate time and resources to investing in a largely entirely new supply chain and assembly process, both to produce an electric motor instead of a combustion engine and to manufacture a vehicle most suited to being propelled by the former over the latter. No internal combustion engine means no exhaust system or fuel intake system, but it has also meant new problems to resolve, most of which have been connected to

the production and use of batteries. With this, direct investments in chip infrastructure appear rare, despite obvious benefits to car manufacturers.

Secondly, that despite the exhaustive computational processing requirements of sensor-equipped, ML-driven autonomous vehicles (Hind, 2023), node shrinkage is not as big a consideration as it is for manufacturers of mobile devices. Whilst car manufacturers have long been attentive to the weight of the vehicles they produce, aiming to save it at every stage of the production and assembly process, automakers aren't under any considerable pressure to reduce the size of their vehicles. On the contrary, cars have never been bigger. Compact SUVs outsold all other types in Europe in 2021, accounting for about 20% of all those sold, with SUVs holding a market share of over 45% in total (Statista, 2022). More, smaller, transistors on every chip are beneficial to all semiconductor customers, but car manufacturers have a multitude of other operational considerations equally worth their time.

Nevertheless, the semiconductor chip shortage did affect car manufacturers, as mentioned above—despite not requiring the same types of chips as found in flagship mobile devices. As SNV contend, the crisis was largely the result of three interconnected aspects: 'overly pessimistic' (Kleinhans, 2021, p. 18) market recovery forecasts by car manufacturers, the prevalence of just-in-time production methods throughout the automotive industry (and associated low inventory levels), and 'simultaneously strong demand for consumer electronics' (Kleinhans, 2021, p. 18) during lockdowns and stay-at-home orders. As SNV further state, chip fabrication ordinarily takes four to six months, and so any submitted order by fabless automotive suppliers such as Bosch or Denso would necessarily entail a wait. Only 5% of TSMC's net revenue in 2022 was generated from automotive customers, compared to 41% from high-performance computing and 39% from smartphones (TSMC, 2022), meaning car companies were at the back of the queue for their chips. In short, the stop/start nature of the global economy during the long-tail of the pandemic has exposed the vulnerability of multiple taken-for-granted aspects of the way automotive firms operate.

Whilst the unprecedented demand for consumer electronics such as video games consoles and exercise bikes might appear unique to a global pandemic, as McKinsey analysis suggests, there remains considerable

overlap between chips required for such consumer devices and contemporary vehicles (Burkacky et al., 2021). Such an overlap, rather than receding, will only continue to grow. As they summarize, cutting-edge nodes (28 nm or smaller) possess a low overlap with respect to 5G (logic, field programmable gate arrays, application-specific integrated circuits) and IoT edge computing (main processing units, memory). But regarding trailing edge nodes (28 nm or larger), there is a high overlap with respect to electrification (discretes, power management, power supply units) and medium overlap with 5G (radio-frequency switches, duplexers, antenna) and IoT edge computing (sensors, microcontrollers, analogue communication).

In this, automotive firms are competing alongside other manufacturers for chip production space in an already high-demand arena (Yoon, 2023). As McKinsey conclude, ‘chip capacity won’t catch up with demand in the short term ... primarily because of the continued increases in volume and sophistication levels of the chips needed to power ... advanced driver-assistance systems and autonomous driving’ (Burkacky et al., 2021, p. n.p.). Longer term, as I’ve tried to outline in this chapter, ‘the auto industry will need to rethink the way it structures contracts for semiconductor-related sourcing’ (Burkacky et al., 2021, p. n.p.). In short, that it might have to reconsider, ‘at least in part’ (Burkacky et al., 2021, p. n.p.) some near sacred and fundamental parts of the contemporary industry, such as just-in-time production and low inventories.

State and supranational level commitments to funding the construction of fabs, in both the EU and the US, in order to decrease reliance on East Asian manufacturers, are a consequence of how important semiconductor chips are to national economies and especially to key industries such as the automotive industry. A physical lack of chips threatens to throttle the ambitions of car manufacturers looking to turn their models into sophisticated, chip-dependent, sensing devices. Beyond questions of technical viability, perceptive accuracy, legal safety, and public acceptance, securing the continued production of semiconductor chips is necessarily vital to the future of autonomous driving.

European Chips Act: (East) German Exceptionalism

This has involved the EU relaxing state aid rules in order to boost the capacity of semiconductor production across member states, even if ‘connected’ vehicles don’t necessarily require the most cutting-edge chips needed for smartphones and tablets. EU state aid rules, laid down in the 1957 Treaty of Rome, prevent EU member states from providing financial incentives to firms or sectors in order to gain competitive advantages (Davies, 2013). The European Chips Act, proposed in February 2022 intended to relax otherwise strictly mandated and ‘cherished’ (Davies, 2023, p. n.p.) state aid rules, to stimulate the production of semiconductor chip fabs within the EU (European Commission 2023a). Such an unprecedented decision suggests that the European Commission recognizes the significant disadvantage the EU will likely (continue to) have in relation to, namely, East Asia, in the future digital economy.

One way of understanding such a proposal is through the notion of exception, as Will Davies (2013) does. The relaxation of state aid rules for the production of chip fabs suggests that this issue is so critical that usual rules should not apply. Indeed, it is arguably the first time that the EU has invoked a relaxation of such rules in order to defend its own ‘technological sovereignty’ (Rikap & Lundvall, 2021, p. 156), suggesting that securing the future ongoing production of semiconductor chips is central to the economic security of the whole EU bloc. As Scott Lavery writes, this constitutes ‘a decisive break from the past’ (Lavery, 2024, p. n.p.) where ‘in the 1990s and 2000s, Washington and Brussels viewed the development of the semiconductor industry as an example of globalization working as intended’ (Lavery, 2024, p. n.p.).

In invoking exception as a special circumstance, then, one must reason that the situation the EU finds itself in is a notably exceptional moment or ‘crisis’. These supply chain issues, in part affected by the global Covid-19 pandemic, have severely affected the future viability of autonomous vehicles. The argument I thus make in this chapter is that such intersecting crises (or polycrisis) offer an external threat to the possibility of autonomous driving, as a parallel to ‘internal’ threats offered by the

lack of technical progress on automation, object-recognition, and motion planning.

One such exception to EU state aid rules applies to the former East Germany, incorporated into the Federal Republic of Germany since 1990 (Davies, 2013). Article 107 (2)(c) of the Treaty on the Functioning of the European Union (TFEU) states that ‘aid granted to the economy of certain areas of the Federal Republic of Germany affected by the division of Germany, in so far as such aid is required in order to compensate for the economic disadvantages caused by that division’ is considered ‘compatible with the internal [EU] market and thus regarded as an *exemption*’ (European Commission, 2015, p. 1, emphasis added).

This is interesting for two reasons. Firstly, of course, that this historic exemption can be considered a plausible reason for why (West) German automotive manufacturers have opened facilities in the former East over the last 30 years, in addition to taking advantage of the collapse of the Soviet Union in opening facilities in the Czech Republic and Slovakia. Here, German automakers have been able to both offshore *and* ‘inshore’ car assembly over the same time period, benefitting from state aid rule exceptions and the neoliberalization of former Soviet satellite states (Pavlínek & Janák, 2007; Pavlínek, 2008).

Then, secondly, that this can also be seen as an additional underlying factor in the decision for battery suppliers and semiconductor firms to open new facilities in the former Eastern states of Brandenburg and Saxony-Anhalt. That is, in addition to commitments made by the EU in the European Chips Act. Accordingly, the re-unification of Germany in 1990 continues to affect the future development of semiconductor chips and the automotive industry in 2024 and beyond, making Eastern Germany in particular an attractive place to open chip fabs and battery production facilities, as well as whole vehicle assembly plants. Intel chief Pat Gelsinger, for instance, on announcing the construction of a new chip fab facility near Magdeburg, Saxony-Anhalt described it as an ‘ideal place’ (Tagesschau, 2022, p. n.p.). It is this situation that has previously been challenged as unfair, for instance, by Europe of Freedom and Direct Democracy (EFDD), the Eurosceptic group in the European Parliament (European Commission, 2015).

As Davies continues, examples of exemptions of state aid include ‘to promote execution of an important project’ or ‘to remedy a serious disturbance in the economy’ (Davies, 2013, p. 44), conditions outlined in the European Chips Act, sparked by the chip crisis. Here, certain ‘positive externalities’ that serve as possible state aid exemptions include ‘high-end’ research and development (R&D) (Davies, 2013, p. 48), such as chip fabs. Due to their incomparable significance for the future economic (technological) sovereignty of the EU, chip fabs represent ‘a particular type of object that resists easy calculation via market-based techniques of valuation’ (Davies, 2013, p. 48). In other words, that semiconductor chip fabs matter so much to the future of innumerable sectors within the EU (consumer electronics, automotive industry, agricultural sector, etc.) that they cannot either be dealt with solely by the EU market nor can be valued accurately enough through or by it.

As the European Commission’s ‘State Aid Action Plan’ outlined, it was committed to analysing instances of ‘market failure’ (European Commission, 2005, p. 6), where state aid rules might be permitted. The question, thus, is precisely what the trigger for the European Chips Act and the conditional relaxation of state aid rules was, in respect to the production of semiconductor chips. Was it that the market as composed was failing to provide customers (car companies, drivers) with well-priced, adequate quality, or easy to acquire products? Or, on the production side, that the market as composed wasn’t stimulating enough competition because, for instance, of the high investment costs associated with semiconductor investment and R&D? Or, on a more expressly political level, that this was unduly increasing reliance of EU businesses on East Asian semiconductor firms, such as TSMC and Samsung?

There is an explicit reference to Article 107 (3)(c) of the TFEU, on allowing the European Commission to ‘approve State aid to facilitate the development of certain economic activities, if the positive effects of such ... outweigh its potential negative impact on trade and competition’ (European Commission, 2023a, p. 3). In particular, that it must offer a so-called incentive effect, be ‘necessary’ as well as demonstrably ‘appropriate’ and ‘proportionate’ (European Commission, 2023a, p. 3) as decided by the European Commission itself. In support of their proposal, the European Commission state that one of the principal conditions would

be that such facilities would have to be ‘first-of a kind’ within the EU, without any ‘equivalent facility’ existing within member states at present (European Commission, 2023a, p. 3).

The conclusion, therefore, regarding the announced Intel facility in Magdeburg—despite the fact the Act had not yet been implemented—would be that the facility would expect to pass this, and any related tests, ordered by the European Commission. This was arguably a somewhat strange position for Intel, Germany, and the European Commission, with Intel undoubtedly receiving assurances that their investment in Germany would not contravene state aid rules in this instance, until the European Chips Act was enacted.⁴ The consequence of such also being that once a ‘one-of a kind’ facility was announced, presumably, no other further facilities like it could be built with state aid, posing questions about how precise a ‘uniqueness test’ in respect to chip fabs would, or could, be.

Conclusion

This chapter has sought to understand how autonomous vehicles ‘secure’ decisions in two ways. Firstly, as ‘auto-*nomos*’ vehicles, they execute decisions like a sovereign actor. Yet, rather than exclusively controlling and exercising authority over a sovereign territory, the autonomous vehicle acts by exclusively controlling and exercising authority over calculative decisions made whilst driving. Following the work of Bratton (2016), Agamben (2005), and Schmitt (1985 [1922]), sovereignty can be understood through an actor’s ability to suspend the normal state of operation, in favour of a ‘state of exception’, in which the actor is neither necessarily outside of or beyond the law, nor wholly inside it either. This ability to decide when the normal rules apply is the essence of the power of a sovereign actor.

Through the setting of operational thresholds—some of which might exceed legal norms in order to avoid other possible riskier situations

⁴The European Chips Act eventually entered into force on September 21, 2023 (European Commission, 2023b).

encountered by the vehicle—the autonomous vehicle has the possibility of deciding when and how an operational state of exception applies. In this, the autonomous vehicle has the capacity to switch between ‘normal’ and ‘exceptional’ modes at will, thanks to the protocological (Galloway, 2004) form its decision-making architecture takes.

Yet, in tune with the messy, unresolved, relationship autonomous vehicles are beginning to engender with other political entities, there are some rather more familiar forms of sovereign action being carried out within this sphere, also. Thus, secondly, the chapter has sought to unpack the effect that emergent political-economic crises have had on automotive supply chains, offering a different sense of ‘securing’ decisions, rooted in establishing or maintaining reliable supply chains of important technological components and systems essential for the delivery of autonomous driving as a future phenomenon. Here, such compound crises manifest as an entangled ‘polycrisis’ (Tooze, 2022c) amounting to a kind of extant, rolling, tumultuous environment in which actors within the automotive industry and the emerging autonomous vehicle industry can only navigate, rather than definitively solve. In times of polycrisis, there are no single ‘fixes’.

Through an analysis of the ‘chipification’ (Forelle, 2022) of automobility and the recent semiconductor ‘chip crisis’, the chapter examined how the latter has been ‘engineered’ by key actors such as chip firms and nation states, to achieve a certain level of supply chain independence. These economic decisions are representative of broader efforts by the EU to cultivate a form of ‘technological sovereignty’ (Rikap & Lundvall, 2021), intending to de-couple supply chains from reliance on other territories such as East Asia. As the final section of the chapter contends, the form such technological sovereignty has taken with the EU has been dependent upon the historical peculiarities of key states within it, such as the exceptional status of the Eastern German states since the re-unification of Germany in 1990. With the exceptional features of the German automotive industry in mind, the future world of autonomous vehicles within the EU will likely follow a more ‘techno-protectionist’ path, offering significant challenges to the US, China, and global models of free trade.

References

- Agamben, G. (2005). *State of exception*. The University of Chicago Press.
- Amable, B., Guillard, E., & Palombarini, S. (2012). Changing French capitalism: Political and systemic crises in France. *Journal of European Public Policy*, 19(8), 1168–1187. <https://doi.org/10.1080/13501763.2012.709011>
- Associated Press. (2011, June 30). Germany votes to end nuclear power by 2022. *The Guardian*. <https://www.theguardian.com/world/2011/jun/30/germany-end-nuclear-power-2022>
- Babic, M. (2020). Let's talk about the interregnum: Gramsci and the crisis of the liberal world order. *International Affairs*, 96(3), 767–786. <https://doi.org/10.1093/ia/iiz254>
- Bratton, B. (2016). *The stack: On software and sovereignty*. MIT Press.
- Burkacky, O., Lingemann, S., & Pototzky, K. (2021). *Coping with the auto-semiconductor shortage: Strategies for success*. McKinsey. Retrieved October 1, 2023, from <https://www.mckinsey.com/industries/automotive-and-assembly/our-insights/coping-with-the-auto-semiconductor-shortage-strategies-for-success>
- Cariad. (2023). Futureproof of hardware. *Cariad*. Retrieved October 1, 2023, from <https://cariad.technology/de/en/solutions/futureproof-hardware.html>
- Carr, A., & Welch, D. (2024, May 1). Will GM regret kicking Apple CarPlay off the dashboard? *Bloomberg*. <https://www.bloomberg.com/news/features/2024-05-01/will-gm-regret-kicking-apple-carplay-off-the-dashboard>
- D'Eramo, M. (2022, July 19). Sinking Germany. *NLR Sidecar*. <https://newleft-review.org/sidecar/posts/sinking-germany>
- Davies, W. (2013). When is a market not a market?: 'Exemption', 'externality' and 'exception' in the case of European state aid rules. *Theory, Culture & Society*, 30(2), 32–59. <https://doi.org/10.1177/0263276412456567>
- Davies, W. (2023, July 13). Pain, no gain. *London Review of Books*. <https://www.lrb.co.uk/the-paper/v45/n14/william-davies/pain-no-gain>
- Destatis. (2022a). The People's Republic of China is again Germany's main trading partner. *Destatis*. Retrieved October 1, 2023, from <https://www.destatis.de/EN/Themes/Economy/Foreign-Trade/trading-partners.html>
- Destatis. (2022b, July 18). Außenhandel: Ranking of Germany's trading partners in foreign trade. Retrieved October 1, 2023, from <https://www.destatis.de/EN/Themes/Economy/Foreign-Trade/Tables/order-rank-germany-trading-partners.pdf>

- Destatis. (2022c). The main German export product: Motor vehicles. *Destatis*. Retrieved October 1, 2023, from <https://www.destatis.de/EN/Themes/Economy/Foreign-Trade/trading-goods.html>
- Durand, C. (2019). In the crisis cockpit. *New Left Review*, 116(117), 201–222. <https://newleftreview.org/issues/ii116/articles/in-the-crisis-cockpit>
- European Commission. (2005). *State aid action plan – Less and better targeted state aid: A roadmap for state aid reform 2005–2009*. EUR-Lex. Retrieved October 1, 2023, from <https://eur-lex.europa.eu/legal-content/EN/ALL/?uri=CELEX%3A52005DC0107>
- European Commission. (2015). *Question for written answer E-012371/2015*. European Parliament. Retrieved October 1, 2023, from https://www.europarl.europa.eu/doceo/document/E-8-2015-012371_EN.html
- European Commission. (2023a). *European Chips Act – Questions and answers*. European Commission. Retrieved October 1, 2023, from https://ec.europa.eu/commission/presscorner/detail/en/qanda_23_4519
- European Commission. (2023b). *European Chips Act*. European Commission. Retrieved October 1, 2023, from https://commission.europa.eu/strategy-and-policy/priorities-2019-2024/europe-fit-digital-age/european-chips-act_en
- European Parliament. (2022, November 3). EU ban on the sale of new petrol and diesel cars from 2035 explained. *European Parliament*. <https://www.europarl.europa.eu/news/en/headlines/economy/20221019STO44572/eu-ban-on-sale-of-new-petrol-and-diesel-cars-from-2035-explained>
- Ewing, J., & Boudette, N. E. (2021, April 23). A tiny part's big ripple: Global chip shortage hobbles the auto industry. *The New York Times*. <https://www.nytimes.com/2021/04/23/business/auto-semiconductors-general-motors-mercedes.html>
- Financial Times. (2023, February 17). Germany's bold policy shift remains a work in progress. *Financial Times*. <https://www.ft.com/content/d1ab1c9e-746f-4122-b86f-0ea95f5c05bc>
- Fine, A. (2021, October 25). Infineon achieves Guinness World Record with VW ID.4 EV. *NGT News*. <https://ngtnews.com/infineon-achieves-guinness-world-record-with-ev-on-longest-journey-in-a-single-country>
- Forelle, M. C. (2022). The material consequences of 'chipification': The case of software-embedded cars. *Big Data & Society*, 9(1), 1–12. <https://doi.org/10.1177/20539517221095429>
- Galloway, A. (2004). *Protocol: How control exists after decentralization*. MIT Press.
- Gekker, A., & Hind, S. (2019). Infrastructural surveillance. *New Media & Society*, 22(8), 1414–1436. <https://doi.org/10.1177/1461444819879426>

- George, P. (2022, December 14). Volkswagen has given a name to its pain, and it is 'software'. *The Verge*. <https://www.theverge.com/2022/12/14/23508088/volkswagen-software-id4-bug-problem-smartphone>
- Gitlin, J. M. (2022, October 24). After complaints, Volkswagen will ditch capacitive steering wheel controls. *Ars Technica*. <https://arstechnica.com/cars/2022/10/after-complaints-volkswagen-will-ditch-capacitive-steering-wheel-controls/>
- Heaven, D. (2019, November 20). The humble material that transformed the world. *BBC Future*. <https://www.bbc.com/future/bespoke/made-on-earth/how-the-chip-changed-everything/>
- Hill, K. (2024, March 11). Automakers are sharing consumers' driving behavior with insurance companies. *The New York Times*. <https://www.nytimes.com/2024/03/11/technology/carmakers-driver-tracking-insurance.html>
- Hille, R., & Sevastopulo, D. (2022, October 24). TSMC: The Taiwanese chip-maker caught up in the tech cold war. *Financial Times*. <https://www.ft.com/content/bae9756a-3bce-4595-b6c9-8082fd735aa0>
- Hind, S. (2021). Dashboard design and the 'datafied' driving experience. *Big Data & Society*, 8(2), 1–14. <https://doi.org/10.1177/205395172111049862>
- Hind, S. (2023). Sensor work: Enabling the interoperation of autonomous vehicles. *Journal of Computer-Mediated Communication*, 28(5), 1–10. <https://doi.org/10.1093/jcmc/zmad014>
- Hind, S., & Gekker, A. (2022). Automotive parasitism: Examining Mobileye's 'car-agnostic' platformisation. *New Media & Society*, 0(0), 1–21. <https://doi.org/10.1177/14614448221104209>
- Hind, S., Kanderske, M., & van der Vlist, F. (2022). Making the car 'platform ready': How big tech is driving the platformization of automobility. *Social Media + Society*, 8(2), 1–13. <https://doi.org/10.1177/20563051221098697>
- Hind, S., & Seitz, T. (2024). Cynical technical practice: From AI to APIs. *Convergence*, 30(1), 29–48. <https://journals.sagepub.com/doi/10.1177/13548565221133248>
- IBM. (2022). *Silicon Germanium chips*. IBM. Retrieved October 1, 2023, from <https://www.ibm.com/history/silicon-germanium-chips>
- Juncker, J.-C. (2016, June 21). Speech by President Jean-Claude Juncker at the Annual General Meeting of the Hellenic Federation of Enterprises (SEV). *European Commission*. https://ec.europa.eu/commission/presscorner/detail/en/SPEECH_16_2293
- Kleinhans, J.-P. (2021). *The lack of semiconductor manufacturing in Europe*. Stiftung Neue Verantwortung [New Responsibility Foundation]. Retrieved October 1, 2023, from https://www.stiftung-nv.de/sites/default/files/eu-semiconductor-manufacturing.april_.2021.pdf

- Korosec, K. (2023, October 24). California DMV immediately suspends Cruise's robotaxi permit. *TechCrunch*. <https://techcrunch.com/2023/10/24/dmv-immediately-suspends-cruises-robotaxi-permit-in-california>
- Lavery, S. (2024, April 26). Europe's core. *NLR Sidecar*. <https://newleftreview.org/sidecar/posts/europes-core>
- Miller, C. (2022). *Chip war: The fight for the world's most critical technology*. Simon & Schuster.
- NATO. (2022, July 4). Finland and Sweden complete NATO accession talks. *NATO*. https://www.nato.int/cps/en/natohq/news_197737.htm
- Parikka, J. (2015). *A geology of media*. University of Minnesota Press.
- Pasquale, F. (2017, December 6). From territorial to functional sovereignty. *Law and Political Economy*. <https://lpeproject.org/blog/from-territorial-to-functional-sovereignty-the-case-of-amazon/>
- Pavlínek, P. (2008). *A successful transformation? Restructuring of the Czech automobile industry*. Physica-Verlag.
- Pavlínek, P., & Janák, J. (2007). Regional restructuring of the Škoda auto supplier network in the Czech Republic. *European Urban and Regional Studies*, 14(2), 133–155. <https://doi.org/10.1177/0969776407076101>
- Phillips, L., & Rozworski, M. (2019). *People's Republic of Walmart: How the World's Biggest Corporations Are Laying the Foundations for Socialism*. Verso.
- Reuters. (2021, December 14). COVID-19 rattles major Chinese manufacturing province. *Reuters*. <https://www.reuters.com/world/china/production-suspensions-chinese-manufacturing-hub-amid-covid-19-outbreak-2021-12-14/>
- Rikap, C., & Lundvall, B.-Å. (2021). *The Digital Innovation Race: Conceptualizing the Emerging New World Order*. Palgrave Macmillan.
- Schmitt, C. (1985 [1922]). *Political theology*. Cambridge, MA: MIT Press.
- Shankland, S. (2018, September 12). iPhone XS A12 Bionic chip is industry-first 7nm CPU. *Cnet*. <https://www.cnet.com/tech/mobile/iphone-xs-a12-bionic-chip-is-industry-first-7nm-cpu/>
- Statista. (2022). *Number of car registrations in Europe in 2021, by segment*. Statista. Retrieved October 1, 2023, from <https://www.statista.com/statistics/697801/passenger-car-sales-in-italy-by-segment/>
- Tagesschau. (2022, March 15). Intel baut Chipfabrik in Magdeburg [Intel builds Chipfab in Magdeburg]. *Tagesschau*. <https://www.tagesschau.de/wirtschaft/unternehmen/intel-magdeburg-101.html>
- Taiwan Semiconductor Manufacturing Company (TSMC). (2022). *TSMC Annual Report 2022 (1)*. TSMC. Retrieved October 1, 2023, from <https://investor.tsmc.com/sites/ir/annual-report/2022/2022%20Annual%20Report-E.pdf>

- Ting-Fang, C., & Li, L. (2022, August 4). The resilience myth: Fatal flaws in the push to secure chip supply chains. *Financial Times*. <https://www.ft.com/content/f76534bf-b501-4cbf-9a46-80be9feb670c>
- Tooze, A. (2018). *Crashed: How a decade of financial crises changed the world*. Allen Lane.
- Tooze, A. (2020, December 10). Europe's response was a total disaster. *Internationale Politik*. <https://internationalepolitik.de/en/europes-response-was-total-disaster>
- Tooze, A. (2022a, January 21). Chartbook #73: Crisis pictures (Krisenbilder) – Mapping the polycrisis. *Chartbook*. <https://adamtooze.com/2022/01/21/chartbook-73-crisis-pictures-krisenbilder-mapping-the-polycrisis/>
- Tooze, A. (2022b, June 24). Chartbook #130: Defining polycrisis – From crisis pictures to the crisis matrix. *Chartbook*. <https://adamtooze.com/2022/06/24/chartbook-130-defining-polycrisis-from-crisis-pictures-to-the-crisis-matrix/>
- Tooze, A. (2022c, October 28). Welcome to the world of polycrisis. *Financial Times*. <https://www.ft.com/content/498398e7-11b1-494b-9cd3-6d669dc3de33>
- Van Loo, R. (2020). The new gatekeepers: Private firms as public enforcers. *Virginia Law Review*, 106, 467–452. https://scholarship.law.bu.edu/faculty_scholarship/800/
- Wacket, M. (2021, November 23). German parties agree on 2030 coal phase-out in coalition talks – sources. *Reuters*. <https://www.reuters.com/business/cop/exclusive-germanys-government-in-waiting-agrees-phase-out-coal-by-2030-sources-2021-11-23/>
- Yeung, H. (2022). *Interconnected worlds: Global electronics and production networks in East Asia*. Stanford University Press.
- Yoon, J. (2023, August 29). Nvidia's critical weak point: Taiwan. *Financial Times*. <https://www.ft.com/content/bec85749-9354-4470-ac88-9323541c7bce>
- Zeitlin, J., Nicoli, F., & Laffan, B. (2019). Introduction: The European Union beyond the polycrisis? Integration and politicization in an age of shifting cleavages. *Journal of European Public Policy*, 26(7), 963–976. <https://doi.org/10.1080/13501763.2019.1619803>



7

Relaxing Decisions: Making Driving Chill

This chapter considers the rise of an open-source project called ‘Comma’. Based in San Diego, US, and founded by George Hotz, the first person to hack an Apple iPhone, Comma has been considered a threat to both big tech dreams of autonomous driving and automotive industry visions of assisted driving. The chapter explores how driving as an emotive, affective, bodily, and cultural experience is considered, by Hotz and Comma supporters, to be under threat by big tech visions of automation. Comma’s offer of an aftermarket, ‘bolt-on’ advanced driver-assistance system (ADAS), based on a smartphone-style device, compatible with over 200 existing vehicle brands and models (Comma, 2022), is thus designed to offer the possibility of a limited form of ‘autonomous’ driving, whilst allowing—so their argument goes—customers to retain in control.

In this I explore how Comma is about ‘making driving chill’, as their tagline contends. Many of their avid customers and supporters suggest on online discussion boards that they desire to achieve what they call ‘that openpilot vibe’, as offered by the company’s ML-driven platform, openpilot. Put simply, Comma can be understood as a San Diego version of what Barbrook and Cameron (1996) called the ‘Californian ideology’, a contemporary manifestation of a belief that crystallized on the west coast of the US during the 1980s and 1990s that ‘computers somehow seemed

poised to bring to life the countercultural dream of empowered individualism, collaborative community, and spiritual communion' (Turner, 2006, p. 2).

To experience the joy of retrofitting their vehicles with a Comma device, users participate in a range of connected activities, implicitly accepting, and often explicitly defending, their 'datafied' driving experience (Hind, 2021), in order to contribute to the improvement of the performance of openpilot. Such activities involve performing various 'hacks' to their vehicles to improve the capabilities of the device itself, routinely discussing and sharing tips on a public—but company-moderated—Discord channel. This avid online user community is arguably unique within the automotive industry, only rivalled by the fierce support Tesla drivers have for their own vehicles, and the company's Full Self-Driving System (FSD). Where Comma differs, as this chapter intends to show, is in their support for open-source ethics, commitment to the 'everyday' driving experience, and a future of autonomous driving where individual drivers are nominally in control.

The chapter begins with detailing the 'empowered individualism' embodied by Comma, before introducing what I call Comma's 'culture of testing'. Through an analysis of the Comma online community, I consider four test phenomena indicative of the 'spiritual communion' felt by Comma users: the calibration of Comma devices, the reporting and tracking of device bugs, the posting of 'vibes', and the cultivation of shared experiences.

The Californian Ideology (Empowered Individualism)

In 1996 Richard Barbrook and Andy Cameron published a polemic on the rise of the 'virtual class' (Barbrook & Cameron, 1996), constituting their now famous critique of the 'Californian ideology'. A mix of 1960s counterculture and a kind of electronic libertarianism, the Californian ideology was Barbrook and Cameron's term for the epistemological foundations for an emerging class of 'technophiliacs' (Barbrook & Cameron,

1996, p. 48). Through a shared understanding of the liberating force of new kinds of digital technology—the desktop computer, the internet—these technophiliacs would come to constitute the driving political dynamic of Silicon Valley from the 1980s onwards. ‘In this version of the Californian Ideology’, as Barbrook and Cameron (1996, p. 53) wrote, ‘each member of the “virtual class” is promised the opportunity to become a successful hi-tech entrepreneur’ (Barbrook & Cameron, 1996, p. 53). New ‘information technologies’ beginning to find a mass market in the mid-1990s were, according to proponents, able to ‘empower the individual, enhance personal freedom, and radically reduce the power of the nation-state’ (Barbrook & Cameron, 1996, p. 53).

It is this discourse of empowerment that Philip Agre (1995) laid out comprehensively, considering not only how new kinds of information technologies were being championed by an emergent virtual class, in opposition to the state, but how they were manifesting changes in corporate work cultures, too. Whilst Agre acknowledges that the notion of empowerment does not wholly emerge from 1960s countercultural ideas around subjectivity and personal agency, he nonetheless suggests that empowerment ‘be understood as a general trend towards the outward and downward delegation of decision-making authority in organizations’ (Agre, 1995, p. 172). In other words, that individual empowerment was emerging as a discourse within the workplace, substantially re-organizing the relationship between ordinary worker and management, resulting in greater ‘operational control’ (Agre, 1995, p. 172) of each workers’ own set of tasks and responsibilities being handed over to workers from management. In such a case, at least in theory, each individual worker would be granted greater autonomy in their day-to-day work, largely free from overbearing managers—either assigned to work on jobs of their own choosing or free to work towards set objectives in any way they (either alone or in teams) saw fit.

Agre’s contribution to this debate around worker empowerment was that ‘decentralized computer technology’ (Agre, 1995, p. 177)—essentially the desktop computer—was increasingly key to how individual empowerment was being engendered on the ground, in workplaces. Even more significantly, Agre came to understand this empowerment discourse as part of a bi-directional process he called the ‘empowerment and

measurement regime' (Agre, 1995, p. 176), a relationship that Agre suggested was rarely considered as part of a 'single, coherent system' in prevailing business literature (Agre, 1995, p. 176). For Agre, the 'empowerment' of the desktop computer (and all its later accoutrements—email, etc.) was only ever made possible by an accompanying process of employee monitoring and accounting. In other words, management weren't either freely giving away their own agency, nor had necessarily been wrestled out of it by the overwhelming social power of the virtual class, but were happily providing it on the basis of some kind of return. Employees would be free to make their own 'localized' decisions, for instance, on how to organize work tasks directed towards a set objective, but in granting this kind of (computerized) agency, workers would have to give up a form of workplace privacy in the shape of ongoing performance management. The empowerment and measurement regime, as Agre thus understood it, would be an 'advanced form of ... synthesis, comprised of one theme drawn from the tradition of corporatism and autonomy – empowerment – and another theme drawn from the tradition of rationalization and control – measurement' (Agre, 1995, p. 180). Decentralized computer technologies, of a new information variety, were becoming fundamental to the entrenchment, and proliferation, of the regime itself.

Comma as Open-Source Project (Collaborative Community)

As Luis Alvarez León has written, the rise of Comma invites a 'fuller political economic examination of self-driving car data' (Alvarez León, 2019, p. 13), in which 'the double-edged sword of empowering users to access and interpret their own vehicle data, while creating a market for it, exemplifies one of the defining tensions of digital capitalism' (Alvarez León, 2019, p. 13). In Comma, open-source ethics—of free access, the right to share software, the right to modify software—routinely come under scrutiny, as the company uses them as the foundation of both its

general business model and the operational foundation of the ‘end-to-end’ data-driven, machine learning approach to autonomous driving.

In this, Comma offers an interesting account of how open-source ethics are used to give credence and social capital to a disruptive tech product, considered by Hotz as an entire new consumer electronic device class (Hotz, 2021). Part of the allure to many Comma users is the possibility to assume the role of an active beta tester—an external, public participant in the development of the Comma device and openpilot software. This is the result both of Comma’s consistent marketing around the aforementioned mantra, ‘make driving chill’, but also a result of legal restrictions imposed on Comma in 2016, by the US National Highway and Traffic Safety Administration (NHTSA), following the release, and subsequent cancellation, of their original device, the Comma One (Etherington, 2016). As a result, all Comma products are sold as developmental devices meant for beta testing only. Although in practice, Comma devices are freely available to purchase, it has bred a significant, active, committed, user community attracted by the freedom it affords users to retrofit their vehicles with autonomous driving-like qualities.

As a result of the above and other aspects to which this chapter will explore, Comma can be considered an adaptation, or more appropriately, a *mutation* of the idea of participatory algorithmic authorship, similar to, and reliant on, other platforms such as Github which operate as a ‘form of distributed and iterative writing in which multiple developers contribute to the rewriting and editing of software’ (Amoore, 2020, p. 97). In the case of Comma, this involves users *actively generating* driving data used to improve the capabilities of the underlying openpilot software, which drives the device and ultimately, the vehicle itself. The difference, perhaps, between Comma and other contemporary examples of data-driven platforms, or platforms utilizing active user ‘use data’, is that this is considered a virtue of the product, rather than a reluctant feature, or necessary evil—a belief that users have ‘full(er) customizable power within, and critically, *of*, the vehicle’ (Hind, 2021, p. 2, authors’ emphasis) as augmented through their Comma device.

Although different in many other ways, one might interpret the enthusiastic Comma user community as comparable with, or divergent from, various other contemporary ‘pioneer communities’ (Hepp, 2019) such as

the quantified self (QS) movement, the maker movement, and hacking communities more broadly. Here, instead of sensing devices being used by individuals to track bodily movements and health (as in the QS movement), a Comma device is used by individuals to track their vehicular movements and health, as extensions of the self, and a constituent of the 'embodied driver-car' (Dant, 2004, p. 71), by way of feeding such data into an automated pipeline designed to improve, refine, and further modulate, the novel experience of driving an autonomous vehicle. Hepp (2019) understands such groups as part of the global spread of 'cybernetic counterculture' (Turner, 2006) emergent from the west coast of the US from the 1960s onwards, discursively rendered as the Californian ideology by Barbrook and Cameron. Here, such movements are indebted to the 'pioneering' practices of new tech communities, taking on a globally mutated form far beyond San Francisco or California itself. What establishes such contemporary communities as continuations of formational movements is their commitment to what Hepp (2019) identifies as four key principles: their self-styled belief to be 'forerunners', their role as 'intermediaries' between different domains and spheres, their engagement in 'experimental practices', and their designing, building, and testing of 'visions of possible future scenarios' (Hepp, 2019, pp. 32–33).

Put in political-economic terms, Comma can be seen as the strange mutation of relations between two camps. On one side, what Barbrook and Cameron (1996) refer to as the 'virtual class', embodying the Californian ideology, or what McKenzie Wark (2004, 2015) has variously referred to as the 'vectorialist class': a class that owns neither property nor the means of production, 'but instead owns the vector along which information is gathered and used' (Wark, 2019, p. 2).¹ This virtual or vectorialist class is directly in conflict with the so-called hacker class (Wark, 2004), and 'as the vectorialist class consolidates its monopoly on the means of realizing the value of intellectual property, it confronts the hacker class more and more as a class antagonist' (Wark, 2004, p. 8). This 'subordinate class' (Wark, 2015, p. n.p.) of workers, the hackers, are

¹ The connection between these two terms, *virtual* and *vectorial*, is their respective Deleuzian qualities.

principally tasked with ‘the production of new information’, as Wark (2015, p. n.p.) suggests:

The production of new information as information is based on a technical separation of the flow of information from its material substrate such that while information still has no existence outside of a material substrate, its relation to that substrate becomes abstract. The potential of this development is then constrained and channeled via elaborations of the private property form. (Wark, 2015, p. n.p.)

In other words, it is the hacker class who are involved in the generation of (new) information that becomes the lifeblood of the vectorialist class through processes of information ‘capture’ (Agre, 1994), to which they must abide by. As Wark pithily writes, ‘information wants to be free but is everywhere in chains’ (Wark, 2015, p. n.p.), much like the hacker class itself, forever bound to the empowerment-measurement regime.

Further, like Wark suggests, ‘the production of intellectual property like the production of anything, requires cooperation and collaboration’ (Wark, 2015, p. n.p.), the kind of cooperation and collaboration facilitated by open-source ethics, and the curation and ongoing mobilization of a user community able, ultimately, to deliver intellectual property in the form of a proprietary device. Wark (2015) characterizes this cooperative/collaborative work as ‘in-sourcing’, where out-sourcing ‘sends a worker’s job overseas to another worker’, in-sourcing ‘assigns the hacker’s job to anyone who will perform the task for free’ (Wark, 2015, p. n.p.). As will be expanded on later in this chapter, there are plenty of willing people who will indeed do this for Comma, for free. That they do so freely and willingly shows just how powerful the vectorialist class have become today and to which those driving the Comma project belong.

Thus, Comma might be understood as fostering a form of ‘cynical technical practice’ (Hind & Seitz, 2024), as open-source ethics have been co-opted by, and folded into, Comma’s commercial enterprise. Through this cynical co-optation, an active user community has been cultivated that displays many of the same characteristics of other existing open-source projects without a commercial dimension. Here, I use the term

cynical technical practice as a response to Phil Agre's articulation of 'critical technical practice' (Agre, 1997).

Likewise, on a practical level, Comma devices act 'parasitically' upon drivers and their vehicles, much like other commercial ADAS (Hind & Gekker, 2024), offering a form of empowerment in exchange for data aggregation (Agre, 1994, 1995). In other words, Comma is not only a contemporary example of the Californian ideology in action but of Agre's empowerment and measurement regime. As subsequent sections will detail, that through the 'inversion' of the firm (Parker et al., 2017), and the externalization of product research and development (R&D), it can be seen as fascinating case of the externalization of the empowerment and measurement regime, such that it is not only 'internal' employees subject to these intertwined forces but 'external' users, developers, and other interested parties.

Antagonists: Coneheads and Consumer Reports

One of the central tenets of the Comma project is broad antagonism towards two camps. Firstly, to big tech-driven autonomous driving, which it sees both on a technical level as overcomplicating the design of the perceptual components of autonomous vehicles and on an experiential level, as ruining the sacred agency of the car owner. Secondly, to automotive manufacturers, which it considers both as resistant to DIY automotive cultures and a laggard with respect to designing cheap, effective, driver-assistance devices. Put differently, Comma desires to disrupt both big tech and automotive industries at large. In this section I want to work through each of these antagonisms in turn, as they provide a good sense of the *raison d'être* of the Comma project itself.

In July 2021 Comma CEO George Hotz hosted the company's first 'COMMA_CON' event, a product launch of their new ADAS device, the Comma Three, 'akin to Facebook's F8 or Google's I/O' (Hind & Seitz, 2024, p. 38). 'If Tesla's the iOS [of self-driving cars], we're going to be the Android', Hotz announced (Hotz, 2021). Continuing the

comparison with Elon Musk's electric vehicle company, Hotz claimed, 'we like to say the Tesla slogan is: "look at this crazy feature"' (Hotz, 2021), a marked difference to Comma's 'make driving chill'. Hotz goes on to ask the attendees—largely a collection of committed Comma users and owners—'what does it mean to be "good"?' (Hotz, 2021). 'What makes a system good isn't the crazy things that it can do', Hotz continues, 'but it's the things that it can do over and over again, repeatedly and well. That's what chill means to me' (Hotz, 2021). In this statement of the Comma project's guiding principle and ethos, Hotz makes repeated reference to Tesla's FSD 'Autopilot' feature. Far from autonomous, Autopilot is best understood as a rival ADAS, able to deliver automated assistance to drivers in a specific range of situations, such as exiting a motorway (Tesla, 2023). As Hotz considers, any fancy demonstration of an ADAS performing more experimental manoeuvres ('crazy things') in a more anxiety-inducing manner is less important than executing certain manoeuvres that render the driving experience relaxing.

The notional figure of this first technical antagonism is the so-called cone guy or conehead, a figure that Hotz frames as the epitome of the wrong-headed approach to the automation of driving led by tech firms and the likes of Tesla (Fig. 7.1).² Moving on from his critique of Tesla,

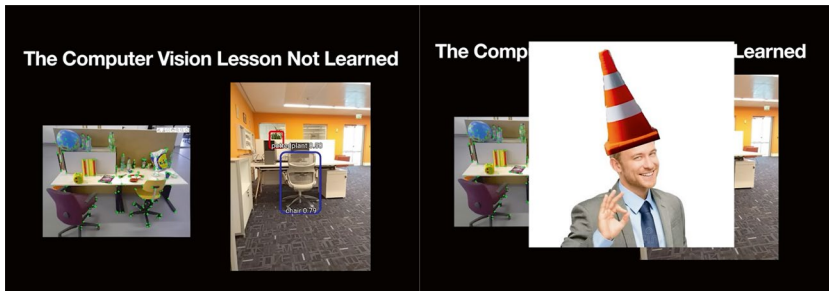


Fig. 7.1 The 'computer vision lesson not learned' and the figure of the 'cone guy'. (Source: Hotz, 2021)

² Despite being global electric vehicle manufacturer with a market cap of \$767 billion (as of August 2023), I prefer to place Tesla alongside other big tech firms, rather than competing automotive manufacturers, because of their software-first approach to automobility, amongst other reasons.

Hotz asks whether ‘these people [can] think a little higher up in the abstraction space?’ (Hotz, 2021) when it comes to delivering automation. In this, Hotz considers what he calls ‘the computer vision lesson not learned’ (Hotz, 2021): that the task of object-recognition and prediction is not the same as the task of driving itself. For Hotz, ‘it’s unclear how you go from prediction to driving’ (Hotz, 2021), and focusing on the former—as computer vision researchers do—still leaves them ill-equipped to solve the task of self-driving:

Why is this lesson still not learned, and why are these companies still hiring cone guys?! You’ve got to get into the taxonomy of cones. You have small cones, you have large cones. What does it mean if a cone is “smushed”? What about the big cones like this [gestures]. The thing outside? Is that a cone?! You gotta hire a cone guy for that, right! So the lesson’s really, err, not learned.’ (Hotz, 2021)

Thus, the ‘cone guy’ encapsulates what Hotz considers as the wrong approach to delivering autonomous driving: an express focus on categorizing objects and of simply building up an ever-larger, ever-more complex ‘prop stash’ (Hind, 2019, p. 412; Madrigal, 2017) of objects liable to be encountered by an autonomous vehicle. While Hotz doesn’t explicitly mention an autonomous vehicle company, it’s clear he’s talking about Waymo, and other firms (at the time, Uber ATG), who take a taxonomical approach to automation, as Chap. 5 contended.

The second antagonist is the automotive manufacturer, as rendered through the lens of consumer rankings. Referring to a consumer report, Hotz ‘mentions how Comma beat all 18 competitors’ (Hind & Seitz, 2024, p. 39) in a comparison between the ADAS capabilities of automotive manufacturers, including Tesla’s Autopilot, Mercedes-Benz’ Driver Assistance, and Volvo’s Pilot Assist (Consumer Reports, 2020). Here Hotz is critical of such manufacturers—many of whom were openly hostile towards Comma—defiantly stating that ‘we did ship a Comma 2, and it did embarrass the car industry’ (Hotz, 2021). Understood as technological laggards, unable to compete with Comma, Hotz sees automotive manufacturers as largely uninterested—or unable—to deliver high-performing driver assistance. Their opposition to Comma on the

basis of safety regulations, in respect to Consumer Reports (2020) rankings noted by Hotz, appears unjustified in light of the overall ratings.

These antagonisms can be seen as evidence of ‘intra-ruling class conflict’ (Wark, 2015, p. n.p.) evident in the autonomous driving domain, each drawing a slightly different set of relations to the subordinated hacker classes. In Wark’s analysis of such conflicts between landlords and capitalists in the work of David Ricardo, ‘the more of the surplus landlords can capture in the form of rent, the less there is for capitalists to capture in the form of profit’ (Wark, 2015, p. n.p.). In this formulation, the landlord class is the passive operator, the ‘octopus’ or parasite extracting wealth merely off the work of others. By contrast, the capitalist must work to re-invest acquired profits or risk losing out altogether to other more committed members of the capitalist class. For Comma, their status in a comparable intra-class battle is one tied most closely to the vectorialist class Wark speaks of, neither linked to rents nor profits, land nor industry, but control over the flow of information and a subordinate class who generate and channel it: the hacker class. In this specific form, one is tempted to put a different, even more appropriate name to it: the Comma class.

Comma Culture or the Comma Class (Spiritual Communion)

The Comma community is central to the project itself: how it is organized, what it prioritized, how devices are tested, and how users foster a sense of belonging with like-minded people. At the centre of the Comma community is Discord: an online message board-style platform. Communities on Discord are organized through ‘servers’ on which users gain access to the communities they desire, with server hosts able to set up different ‘channels’ ordinarily corresponding to different topics. On the ‘Comma.ai community’ server, first launched in July 2020, users are presented with eight different channels, including categories to help users orientate themselves (Onboarding), discuss general Comma-related issues (General), provide feedback on the Comma user experience

(Feedback), and participate in vehicle model-specific conversations (Vehicle-specific). A fairly stringent set of community rules dictate the type and location of posts that should be made by users, driven by a guiding ethos: the ‘comma.ai community is centered around openpilot and improving the driving experience’ (VirtuallyChris, 2020), and a reminder that ‘if the company wins we all win’ (VirtuallyChris, 2020). Any users who see things differently are kindly asked to ‘leave this private discord’ (VirtuallyChris, 2020).

By autumn 2023, the Comma Discord channel had 30,770 members, up from 26,448 in spring 2022. On entering the channel one afternoon in late summer, 3163 members were currently online. One person who counts himself amongst the active userbase is George Hotz himself, who routinely posts updates of Comma projects, as well as responding to—and calling out—users with specific gripes or issues with Comma or Comma products. Whilst much of the conversation taking place on the Discord server is typical of online fora, displaying all the characteristic style and tone of online message boards, it is home to an extremely active user base committed to ‘improving the driving experience’, as the pinned post in the Onboarding channel suggests, passionately dedicated to ‘making driving chill’ and ‘chasing that openpilot vibe’ in reference to the thrill users get from using Comma’s openpilot software.

Before breaking down these aspects of what I call Comma’s *culture of testing*, it is worth situating it within a broader context of user testing and the testing of digital technologies. Firstly, Comma can be seen as an example of the far longer history of DIY tech culture, typified by the free and open-source (FOSS) movement. In such cases, users are encouraged to gain and share skills relating to the upkeep of different technologies and devices, driven by collective opposition to proprietary systems and hardware, and digital rights management (DRM)-style approaches to digital content and software (Stallman, 2015; Tkacz, 2015). Comma users—despite the ‘cynical’ nature of the operation—certainly align with a version of these traditions, albeit refracted through Comma company ethics. As such, Comma seems to ‘echo the utopians of the 1990s’ (Turner, 2006, p. 33), framing the project as reimagining ‘the rebirth in hardware and software of a single, “free” culture ... outside the mainstream’ (Turner, 2006, p. 33), as well as overtly antagonistic towards it.

Secondly, Comma's culture of testing can be seen as an example of the 'inverted firm' (Parker et al., 2017) in which companies outsource R&D to external parties, such as developers. In this case, whilst this outsourcing is not exactly wholesale (Comma does still have an internal development team), they nevertheless rely on an active, knowledgeable user base to inform internal teams on emerging technical issues or bugs—something I will discuss in more detail later. The Comma user community is not only involved in testing Comma devices but also scoping future compatibility possibilities (i.e. with new car models), and even modifying and 'forking' them for enhanced performance. Of further interest is that the practice of inverting the firm in such a manner, and of outsourcing certain aspects of R&D to users, is novel to automotive manufacturers themselves. Here, automotive manufacturers have historically preferred to outsource certain kinds of R&D to trusted suppliers, under the watchful eye of Original Equipment Manufacturers (OEMs), ordinarily dictating the work performed by such suppliers (Hind et al., 2022; Pavlínek & Janák, 2007).

In contrast, this externalization of testing, research, and product development is seen as a key operating principle of digital platforms, and their fostering of external relations between the firm and wider developer 'ecosystems' (van der Vlist, 2022; van der Vlist et al., 2022). Even more specifically, some firms within the domain of autonomous driving (Waymo, Argo AI) have designed particular mechanisms, such as the hosting of (external) annual 'challenges', to contribute to the (internal) development of autonomous vehicle systems (Hind et al., 2024).

But what exactly is constitutive of 'Comma culture' or the 'Comma class'? I want to address this question in four parts, proceeding with the kind of necessary 'operational' approach I outline in the introduction to the book. In short: Comma's culture of testing—whilst very much a standalone feature of the community itself—is nonetheless routinely mobilized in the service of 'making driving chill' and achieving, for users, a so-called openpilot vibe. These *test phenomena* concern four connected practices: the calibration of devices, the tracking/reporting of bugs, the posting of vibes, and, ultimately, the cultivation of shared experiences.

Test Phenomenon I: Calibrating Devices

Comma devices are typically referred to as ‘devkits’³—the result of a special order from the NHTSA banning the original release (Etherington, 2016). All future versions of the device (Comma Two, Comma Three) have been shipped as devkits, with every version of the openpilot software denoting their beta status, e.g. 0.9.2, 0.9.3, 0.9.4 (see Fig. 7.2). At present the device is sold freely—but with significant disclaimers regarding their developmental status and therefore inherent risk in use. Offering Automated Lane Centering (ALC) and Adaptive Cruise Control (ACC), Comma devices afford both ‘latitudinal’ and ‘longitudinal’ automated control, meaning that ‘openpilot can accelerate, brake automatically for other vehicles, and steer to follow the road/lane’ (Comma, 2023a, p. n.p.).

Ensuring the devices work correctly involves calibration. Their status as devkits only emphasizes the need for them to be set up for local conditions. The principal component of this localization process concerns the mounting of the device on the inside of a vehicle windscreen. If mounted improperly, the owner risks the Comma device being unable to sense the




Fig. 7.2 Latest release notes for Comma openpilot. (Source: Comma, 2023a)

³ Short for ‘development kit’.

road accurately, and amongst other things, determine the speed and angle the vehicle assumes.

Yet, mounting a Comma device isn't necessarily straightforward. One user encountered on the Comma Discord provides help in multiple forms: a short step-by-step written guide with pictures on the blogging site Medium, and a link to a YouTube video narrating the same calibration process (Fig. 7.3). The narrated video is typical of the Comma community: a user sits in their vehicle, ostensibly parked in a home garage with a typical array of garage items in the background (washing basket, boxes, shelves). In the foreground, a mounted Comma device (screen on), another mobile device (screen off), and the vehicle's entertainment system visible below. In the mid-ground the most important object in this calibration process: a tape measure repurposed as a DIY 'plumb bob' used to measure the centre of the vehicle windscreen. In the photos posted in the guide, the reader sees behind the scenes: a metal chain ensuring the tape measure hangs straight, and two plastic water bottles placed on the vehicle hood, cleverly used to extend the plumb line towards the windscreen (Fig. 7.4).

Nothing about this calibration process screams ideal test conditions, but everything—the repurposed plumb bob, the garage packed with



Fig. 7.3 Calibration, Comma community style. (Source: Eyezenheim, 2022b)

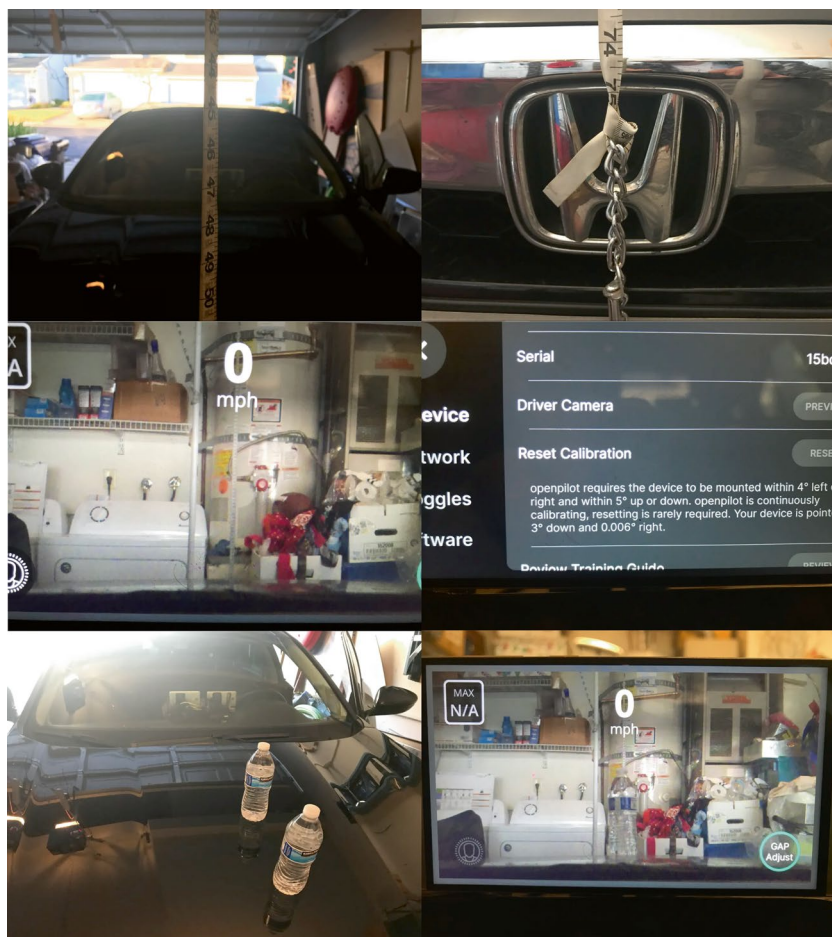


Fig. 7.4 Calibration, behind the scenes. (Source: Eyezenheim, 2022a)

miscellaneous objects, the water bottles; but also the Medium post and the narrated YouTube video—typifies the Comma approach to calibration. Here, the localization process is turned almost hyper-local, executable with an array of ordinary objects close to hand—both general and everyday in type (tape measure, plastic water bottle), but necessarily taken from that most specific and unique of places: the domestic garage.

Earlier, I made reference to the so-called prop stashes used to test autonomous vehicle's ability to perceive miscellaneous road objects. Here another prop stash is also present but in a slightly different form. The domestic garage figures as an even lower-tech prop stash of sorts, integral in ensuring the perceptive capabilities of the Comma device are properly determined by their users. Whilst they have no role here in actually 'testing' these capabilities, they play a significant part in calibrating them for such testing out on the road itself. Whilst these miscellaneous objects are positioned in front of the device, theoretically subject to its gaze, they do not find themselves incorporated into it. Neither object—the tape measure or the water bottles—is subject to machinic object-recognition. The tape measure might assume a different purpose, shifting from measurement device to an instrument used to determine a vertical datum, but at no point does it do so as a prop to be sensed or categorized, but as a tool for subsequent real-world 'props' to be sensed. The water bottles might be elevated to the status of a calibration instrument rather than a humble vessel, but still avoid being subject to the calculative impulses of an ML model.

Test Phenomenon II: Reporting/Tracking Bugs

Once the individual has calibrated their device and put it to use, they might well encounter subsequent issues that need addressing. Some of these issues might be resolved quicker with help from others, including the Comma team themselves. Some might be shared problems—in which case, alerting others would help determine their prevalence. As such, Comma has a dedicated process for 'reporting bugs' involving users raising 'tickets' on a dedicated Github page, and offering a 'Feedback' channel on their Discord server for users to share feedback on different aspects of the Comma experience, from driving routes (#driving-feedback) and driver monitoring (#driver-monitoring-feedback), to openpilot (#nav-feedback) and the device's connected features and subscription service (#connect-feedback).

Despite this, users are typically discouraged from reporting bugs altogether—at least until they've done a few things. For issues with openpilot

and Comma devices, specifically, users have a six-point checklist. Firstly, they have to ensure the bug truly affects their Comma device—rather than their make or model of car. Secondly, if the bug relates to driver monitoring capabilities, users are re-directed to the feedback channels. Thirdly, users must be running the latest version of the openpilot software. Fourth, they must be using officially supported hardware. Fifth, they must have checked the bug doesn't exist already (to which they should check the 'bug tracker' if so). Then, finally, they must ensure they are running the 'stock' version of openpilot (Comma, 2023b). Only then are users given the green light to report a fresh, new bug. If such a spiritual communion exists between Comma users, it's through the diligent adherence to bug reporting that it is affirmed (Figs. 7.5 and 7.6).

One user demonstrates this process in action, asking 'is Comma aware of the steering issue in [version] 8.13?' before detailing a hairy moment the user experienced: 'while talking a turn with openpilot engaged, the steering wheel swings to max turning radius and locks there flashing the "take control immediately" message'. To support their claim, the user ends by stating 'I thought it was something I did last night until someone mentioned it on Reddit this morning'. On posting their message to the #openpilot-experience channel, another user intervenes to remind the poster of the correct protocol: 'Comma does not actively monitor Reddit

The screenshot shows the GitHub 'Issue: Bug report' form for the 'commaai / openpilot' repository. The form is titled 'Issue: Bug report' and includes a subtitle: 'For issues with running openpilot on your comma device. If this doesn't look right, choose a different type.' The form contains several sections: a 'Title' field, a 'Before creating a bug report, please check the following' checklist, a 'Describe the bug' section with a text area, a 'Provide a route where the issue occurs' section with a text area, and an 'openpilot version' field. A sidebar on the right shows navigation links: 'Issues', 'Pull requests', 'Discussions', 'Actions', 'Projects', 'Wiki', 'Security', and 'Insights'. The 'Issues' link is highlighted.

Fig. 7.5 Bug report. (Source: Comma, 2023b)

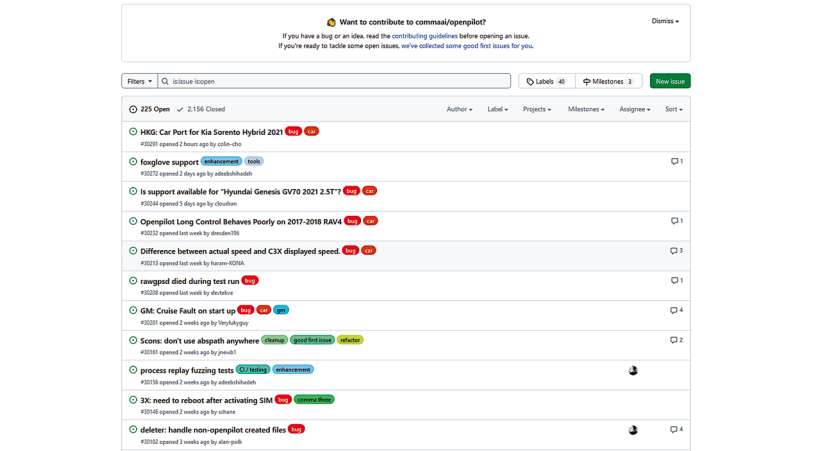


Fig. 7.6 Bug tracker. (Source: Comma, 2023c)

or Discord for bugs and issues. Best is to submit a bug report on their GitHub to grab their attention'. No ordinary member, the other user holds a number of roles that validate their intervention, according to their Discord profile: openpilot contributor, Hyundai/Kia/Genesis moderator, dev (developer), car tester, and fork maintainer.

Upon acknowledging the clarification from the esteemed user, the user with the worrying steering problem mentions they first 'wanted to know if it was a known issue' before raising a ticket. Sure enough, a short while later, the user shares their bug report—'steering bug during turns'—logged on the requisite GitHub page.

What follows is a fascinating account of how the bug is resolved. As the user describes in more detail than their original post:

During a right turn on my commute on 2/22 5:10-ish pm with open pilot [sic] engaged, the car's steering wheel flipped from about 75 degrees to 105 degrees and locked there. The screen flashed orange to take control. I dis-engaged OP [openpilot] (I think, it happened fast) reset the driving line until the lane turned straight and re-engaged cruise. I initially thought it was a one off issue until someone described the exact same problem on the Comma subreddit this morning.

To make absolutely sure, the user posts both a video and a route of the incident occurring, taken from Comma Connect.⁴ On the same day, a Comma employee responds, posting a graph of the user's driving data from the incident, showing the angle of the steering wheel, alongside a graph representing steering torque and a final graph showing engagement of openpilot. Through comparison of these graphs, the Comma employee deduces the user started to take control of the steering wheel when it was at a 40-degree angle—initiating an override of openpilot as the driver proceeds to take the turn. Asking the user to verify where precisely the incident occurred on the graphs posted ('you can select the section by dragging on the timeline, then copy the URL in your address bar') without subsequent reply, the Comma employee offers what is to be their final observation: 'the take control alert was caused by an internal angle limit in your car, Hyundai limits you to ~90 degrees and usually won't steer past that'.

Three graphs seem to confirm the interpretation: steering angle is recorded as 94.5 degrees, a steer warning activates at exactly the same moment, and cruise control is enabled (i.e. cancelled) at the same time also (Figs. 7.7 and 7.8). Whilst the original user appears to go missing, the Comma employee seems to have cracked the case, and the ticket is closed two days after the original report—one of over 2000 tickets that have reached some kind of resolution. Far from unique, this example is the perfect encapsulation of collaborative ethos of the Comma project—a combination of different levels of users ('ordinary' community members, openpilot contributors), a dedicated online community platform for asking questions (Discord), a seemingly parallel community on another (Reddit), a standardized process for logging specific issues (bug report), a platform on which such issues are assessed (GitHub), and another platform that allows both reporter and resolver to visually evaluate associated driving data and, ideally, identify or fix the issue itself (Comma Connect). If, once again, such a spiritual communion exists between Comma users, then it is through this interconnection of protocols and platforms that it is sustained.

⁴ Comma's online device data platform.




Fig. 7.7 Steering bug data. (Source: GitHub, 2022)

Test Phenomenon III: Posting Vibes

Yet, this is a community that rarely is ever strictly engaged in ‘technical work’. Amidst the calibration of devices, the raising of tickets, and the interpretation of driving data, users film, edit, and—most importantly—post content of themselves using their Comma devices. Here the sharing of such media is not somehow outside of or beyond such necessary work but part-and-parcel of being a member of the Comma community. ‘Posting vibes’ is just as non-negotiable as any other activity within the community, equally as vital to the spiritual communion as anything else. Sharing one’s use of the device is—alongside using it—necessarily part of the fun, integral to ‘making driving chill’.

One user typifies this, having posted over 900 videos to their YouTube channel dedicated to sharing their openpilot experiences. Many of the videos are cross-posted on relevant Discord channels, some of which can




Fig. 7.8 Steering angle limit detected. (Source: GitHub, 2022)

be described as technical videos designed to assist other users with calibrating and fixing their devices. However, many arguably serve a different purpose, providing subscribers (they have over 700 on YouTube) and other users and viewers with a sense of how driving with openpilot *feels*. Most videos are descriptively titled, denoting the vehicle and route taken of the user. But some are styled as ASMR videos, those designed, arguably, to offer heightened sensory responses for viewers (Lopez, 2018).⁵ Others make use of Comma's oft-repeated tagline, 'make driving chill', and 'that openpilot vibe' of driving non-handed. In these latter videos there are no step-by-step 'how to' guides. Nor do they offer narration to the viewer on what is happening in the videos.

Some videos elevate the viewing experience to a higher sensory level with breezy instrumental music playing over the top—illustrative of the 'lo-fi' genre that has become the backdrop to online content in recent

⁵ ASMR stands for 'autonomous sensory meridian response', now denoting a popular video genre purportedly designed to stimulate a desirable sensory reaction.

years (Alemoru, 2018; Zaraczynski, 2020). Indeed, such videos can easily be seen as extensions of the ‘slow TV’ genre that has spread through linear programming and non-linear online content centring base pleasures around rhythm and routine—as a stated antidote to the chaos of modern life. For Comma users, ‘making driving chill’ is an optimal, desirable, ideally embodied experience—the videos go some way to showcasing, if not altogether tempting, that experience in lieu of actively experiencing it.

One such video is indicative: just over four minutes long, comprising of a Comma-aided ‘hands-free’ drive along a US highway. The video maintains a fixed position in front of the driver, looking out at eye level through the windscreen, onto the road ahead. Only the driver’s left hand is visible, placed calmly on their knee. The Comma device is clearly in shot, positioned to the top right of the left-hand side steering wheel, mounted on the plastic dashboard of the Honda. From passing road signs the viewer is able to deduce the location of the video: heading away from San Francisco, making their way up the California coastline on the Interstate 80 (I-80).

What elevates the video is the choice of music soundtracking this otherwise archetypal example of a snippet of slow TV: the song ‘Dreams’, by Fleetwood Mac. In October 2020, TikTok user Nathan Apodaca uploaded a video of himself longboarding down a highway in the US state of Idaho after his car had broken down (Apodaca, 2020). Sipping from a cranberry juice carton in his left hand, Apodaca is seen miming to a song: Dreams (2004 remaster) by Fleetwood Mac. Over 92m views, 14m likes, and 710,000 shares later, the video conveyed a sense of bliss—at a moment (during the COVID-19 pandemic) when few were carefree (Beaumont-Thomas, 2020). As the New York Times reported, Apodaca remarked that it was ‘just a video ... that everyone felt a vibe with’, happy he could just ‘chill the world out for a minute’ (Morales, 2020).

The Comma video, then, is an obvious attempt to emulate the carefree bliss of Apodaca’s video. That both take place on US highways/freeways is arguably no coincidence. Whilst Apodaca might be using one to escape his vehicular problems—and by extension, his own worries and stresses—the driver in the other video is using their Comma device to escape theirs. In both, the ordinary motor vehicle is depicted as the problem: for Apodaca using public donations following the viral video to ‘spend on

vehicle repairs and upgrades' (Beaumont-Thomas, 2020, p. n.p.). For the Comma driver, the monotony and boredom of the regular highway commute. The solution for Apodaca turned out to be a 'cranberry red Nissan pickup' (Morales, 2020, p. n.p.) donated by Ocean Spray (sensing an obvious PR opportunity). For the unnamed Honda driver on the I-80, the Comma device offered an emulation of that Apodaca-cum-Fleetwood Mac vibe.

Test Phenomenon IV: Cultivating Shared Experiences

Ultimately, after this calibration work, bug reporting/tracking, and vibe posting, Comma users are in it for one thing, without which they wouldn't be part of the community at all. Owning a Comma device is about *cultivating shared experiences*, a utopian, personalized existence. In short, about making driving chill, matching comfort with style, control with release. The calibration of Comma devices, the fine-tuning of them through the sharing and reporting of bugs and errors, and the posting of driving content are all in service of cultivating and then amplifying that affective, shared experience of 'driving' a car autonomously. Generating an affective 'autonomous' experience whilst retaining personalized power and control is a marked difference from any experience offered either by established car manufacturers or big tech companies. At least, that is, in the eyes of Comma and their avid community of users.

The perfect embodiment of this aesthetic is the #openpilot-experience channel on the Comma Discord server: 'a place for posting openpilot videos and experiences', as the last section demonstrated. In addition, however, besides this invocation to post is the added detail that 'slow mode is on' to encourage 'high quality posts' from users. Discord describes slow mode as 'the most convenient way to make your channel chill out', limiting 'the number of messages a user is able to send in a channel based on a timed cooldown' (Discord, 2021, p. n.p.). Once again, the notion of 'chill'—posting, doing, acting—circulates, imploring users in this specific channel to think and post differently. Whilst the Comma experience is generally oriented towards making *driving* chill, with other Discord channels built to service this end goal, these other channels are not

expressly ‘chill’ in themselves—quite the opposite. The #openpilot-experience channel thereby offers a subsequent deepening, or doubling, of the ‘chill vibes’ that otherwise are meant to pervade whilst driving. Alongside ‘make driving chill’ stands ‘make posting about driving chill’, too.

Another way of understanding this deepening/doubling of the chill experience is that the desired driving experience here is matched by a desired user experience (UX) of the Discord channel. Combined, these constitute an overall, desired *customer experience* to which both Comma device and Discord channel are in service of. As Pine II and Gilmore (2011, p. 23, authors’ emphasis) contend, ‘manufacturers ... must explicitly design their goods to enhance the user’s experience – essentially *experientializing* the goods – even when customers pursue less-adventurous activities’. Here, the manufacturer must always pay acute attention to the use and use context of the goods/service they’ve designed, in order to craft and cultivate the experience of using them. Indeed, Pine II and Gilmore (2011, p. 23, authors’ emphasis) suggest that ‘automakers do this when they focus on enhancing the *driving* experience’, perhaps the most obvious case of needing to attend to the experience of product usage, ‘but they must also focus on other non-driving experiences that occur in cars too’ (Pine II & Gilmore, 2011, p. 23).

For Comma—not a car manufacturer, of course—this is also definitively obvious. There is, evidently, a looping or feedback mechanism at play with Comma, connecting the use of a Comma device (the *driving* experience) to the discursive representation of using the Comma device (the *posting* experience). Whilst the driving experience—using a Comma device to drive autonomously—is ordinarily the end goal, the #openpilot-experience channel crystallizes and supports it, offering the opportunity for users to share a *collective* user experience.

As many have considered, the experience of ‘comfort’ plays a continued role in selling and sustaining private car consumption (Kent, 2015; Sheller, 2004), ordinarily accompanied by the feeling of ‘action’ and ‘effortlessness’ (Kent, 2015, p. 735) whilst driving a car. More specifically, as Sheller writes:

Car consumption is never simply about rational economic choices, but is as much about aesthetic, emotional and sensory responses to driving, as well as patterns of kinship, sociability, habitation and work. (Sheller, 2004, p. 222)

Thus,

We can ask how feelings for, of and within cars occur as embodied sensibilities that are socially and culturally embedded in familial and social practices of car use, and the circulations and displacements performed by cars, roads and drivers. (Sheller, 2004, p. 222)

For Comma users, these ideas of the ‘sensory responses to driving’ and ‘patterns of kinship’ and sociability are decidedly true. This looping of driving experience with posting experience is necessarily attuned to these affective connections. However, Comma users do not simply strive for the same kind of experience as other car drivers. Enabling a distinction between other car drivers and themselves is integral to this ‘unique’ Comma experience—and what subsequently binds the community. This is built firstly on a belief that car drivers in general necessarily yearn for the kinds of comfort, effortlessness, and freedom typically sold to them by car manufacturers—but which cannot be delivered by actually participating in the act(s) of driving itself. Nor, of course, can this feeling be cultivated by the likes of Google/Alphabet and other big tech autonomous vehicle firms who force drivers to cede total control of ‘their’ vehicles to the machine.

Thus, what binds the Comma community, as hinted at previously, is the ‘spiritual communion’ achieved through testing, tweaking, and hacking their Comma devices and—by extension—their beloved automobiles. In so doing, they chase a level of comfort, effortlessness, control, and freedom originally promised by car manufacturers, for which all other drivers yearn but cannot attain. In ‘hacking’ the driving experience, Comma users believe they have elevated themselves onto another plane of existence, embodying an otherwise elusive future of driving few others have ever experienced.

Conclusion

In light of everything, Comma is an anomaly. Neither an automotive manufacturer nor a big tech firm, nor funded by either, Comma is already somewhat different from the competition. Neither committed to building an autonomous vehicle system nor a typical ADAS, it instead is committed to a project somewhere in between. Neither able to leverage huge computational resources nor established engineering know-how, Comma has still managed to launch a standalone ADAS nominally able to satisfy the autonomous driving urges of avid customers.

Amongst all these anomalies—however significant—what marks Comma as unique is their dogged commitment to the Californian ideology (Barbrook & Cameron, 1996), that special brand of electronic libertarianism once the preserve of nearly all Californian tech firms. What this chapter has sought to do, then, is document this ideology in action: the spirit of which flows through manifold aspects of their operation. The protagonists are all of them, Hotz of course, but many more collectively. The antagonists are multiple too: consumer electronics critics, Waymo, Tesla, the NHTSA, and, naturally, the cone guys. Precious few other operators within this space marshal support quite like Comma, and fewer still mobilize such a long line of enemies in order to further their own cause. The spiritual communion of the Comma class is strong, perhaps unbreakable.

The reason why Comma is such a fascinating case study of the development of autonomous driving is because it surfaces an undeniable truth that neither automotive manufacturer nor big tech firm can answer. Drivers—many more than either of them realizes—want automation but on their own terms. Car manufacturers, fighting for new revenue streams, and constantly battling to lower costs, increasingly cannot provide what their customers ultimately desire. Big tech firms, hypnotized by the dream of monopoly control, have little interest in catering to real or imagined drivers who *themselves* yearn for ultimate control of how, and whether, to automate their own beloved vehicle.

Making sense of Comma's 'culture of testing' has been a route to understanding this curious phenomenon. From the calibration of

Comma devices to the reporting of bugs, and from the posting of ‘vibes’ to the cultivation of shared experiences, the project itself occupies a full gamut of explanations for why the Comma project nominally works, drawing people into a mix of technical, social, and aesthetic dimensions.

Chasing what users call ‘that openpilot vibe’, responding to the call to ‘make driving chill’, Comma is fuelled by a forgotten energy at the heart of the Californian tech industry, jettisoned perhaps when Google dropped its unofficial ‘don’t be evil’ motto (Conger, 2018). This chapter has not intended to call for a return to such an ethic, nor for a renewed commitment to electronic libertarianism, but to recognize and examine how the Californian ideology has found a kind of afterlife, living on amongst the ruins of platform capitalism. Far from diametrically opposed to the latter, however, the Comma project has nonetheless found a way to position itself against it, in part, through a somewhat cynical use of technical and cultural critique (Hind & Seitz, 2024). They may not be the future of autonomous driving, but they have certainly altered the present search for it.

References

- Agre, P. (1994). Surveillance and capture: Two models of privacy. *The Information Society*, 10(2), 101–127. <https://doi.org/10.1080/01972243.1994.9960162>
- Agre, P. (1995). From high tech to human tech: Empowerment, measurement, and social studies of computing. *Computer Supported Cooperative Work (CSCW)*, 3, 167–195. <https://doi.org/10.1007/BF00773446>
- Agre, P. (1997). Toward a critical technical practice: Lessons learned in trying to reform AI. In G. C. Bowker, S. L. Star, W. Turner, & L. Gasser (Eds.), *Bridging the great divide: Social science, technical systems, and cooperative work* (pp. 131–158). Lawrence Erlbaum Associates.
- Alemoru, K. (2018, June 14). Inside YouTube’s calming ‘lofi hip hop radio to relax/study to’ community. *Dazed*. <https://www.dazeddigital.com/music/article/40366/1/youtube-lo-fi-hip-hop-study-relax-24-7-livestream-scene>
- Alvarez León, L. F. (2019). Counter-mapping the spaces of autonomous driving. *Cartographic Perspectives*, 92, 10–23. <https://doi.org/10.14714/CP92.1450>
- Amoore, L. (2020). *Cloud ethics: Algorithms and the attributes of ourselves and others*. Duke University Press. <https://doi.org/10.2307/j.ctv11g97wm>

- Apodaca, N. (2020, October 25). Dreams (2004 remaster) – Fleetwood Mac. *TikTok*. <https://www.tiktok.com/@420doggface208/video/6876424179084709126>
- Barbrook, R., & Cameron, A. (1996). The Californian ideology. *Science as Culture*, 6(1), 44–72. <https://doi.org/10.1080/09505439609526455>
- Beaumont-Thomas, B. (2020, October 6). Fleetwood Mac’s Dreams breaks streaming records after viral TikTok video. *The Guardian*. <https://www.theguardian.com/music/2020/oct/06/fleetwood-mac-dreams-breaks-streaming-records-after-viral-tiktok-video>
- Comma. (2022). *Comma Three: Make driving chill with our beautiful new hardware*. Comma. Retrieved October 1, 2023, from <https://comma.ai/>
- Comma. (2023a). *Openpilot*. Comma. Retrieved October 1, 2023, from <https://www.comma.ai/openpilot>
- Comma. (2023b). *Issue: Bug report*. Github. Retrieved October 1, 2023, from https://github.com/commaai/openpilot/issues/new?assignees=&labels=bug&projects=&template=bug_report.yml
- Comma. (2023c). *Openpilot: Issues*. Github. Retrieved October 1, 2023, from <https://github.com/commaai/openpilot/issues>
- Conger, K. (2018, May 18). Google removes ‘don’t be evil’ clause from its code of conduct. *Gizmodo*. <https://gizmodo.com/google-removes-nearly-all-mentions-of-dont-be-evil-from-1826153393>
- Consumer Reports. (2020). *Active Driver Assistance Systems: Test results and design recommendations*. Consumer Reports. Retrieved October 1, 2023 from <https://data.consumerreports.org/wp-content/uploads/2020/11/consumer-reports-active-driving-assistance-systems-november-16-2020.pdf>
- Dant, T. (2004). The driver-car. *Theory, Culture & Society*, 21(4–5), 61–79. <https://doi.org/10.1177/0263276404046061>
- Discord. (2021, October 28). Slowmode – Slllooowwwiiinnng down your channel. *Discord*. <https://support.discord.com/hc/en-us/articles/360016150952-Slowmode-Slllooowwwiiinnng-down-your-channel>
- Etherington, M. (2016, October 28). Comma.ai cancels the Comma One following NHTSA letter. *TechCrunch*. <https://techcrunch.com/2016/10/28/comma-ai-cancels-the-comma-one-following-nhtsa-letter/>
- Eyezenheim. (2022a, January 14). Mount your Comma device using display as guide. *Medium*. <https://eyezenheim.medium.com/mount-your-comma-device-using-display-as-guide-3c270b63ed56>
- Eyezenheim. (2022b, January 17). Centered camera hack openpilot on the rails feeling. *You Tube*. <https://www.youtube.com/watch?v=bEhuJwm0qtk&t>

- GitHub. (2022). Steering bug during turns #23839. *GitHub*. <https://github.com/commaai/openpilot/issues/23839>
- Hepp, A. (2019). *Deep mediatization*. Routledge.
- Hind, S. (2019). Digital navigation and the driving-machine: Supervision, calculation, optimization, and recognition. *Mobilities*, 14(4), 401–417. <https://doi.org/10.1080/17450101.2019.1569581>
- Hind, S. (2021). Dashboard design and the ‘datafied’ driving experience. *Big Data & Society*, 8(2), 1–14. <https://doi.org/10.1177/20539517211049862>
- Hind, S., & Gekker, A. (2024). Automotive parasitism: Examining Mobileye’s ‘car-agnostic’ platformisation. *New Media & Society*, 26(7), 3707–3727. <https://doi.org/10.1177/14614448221104209>
- Hind, S., Kanderske, M., & van der Vlist, F. (2022). Making the car ‘platform ready’: How big tech is driving the platformization of automobility. *Social Media + Society*, 8(2), 1–13. <https://journals.sagepub.com/doi/10.1177/20563051221098697>
- Hind, S., & Seitz, T. (2024). Cynical technical practice: From AI to APIs. *Convergence*, 30(1), 29–48. <https://doi.org/10.1177/13548565221133248>
- Hind, S., van der Vlist, F., & Kanderske, M. (2024). Challenges as catalysts: How Waymo’s Open Dataset Challenges Shape AI Development. *AI & Society*, 0(0), 1–17. <https://doi.org/10.1007/s00146-024-01927-x>
- Hotz, G. (2021, August 4). George Hotz announcing COMMA THREE at COMMA_CON. *YouTube*. <https://www.youtube.com/watch?v=hbLiehrC2DQ>
- Kent, J. L. (2015). Still feeling the car – the role of comfort in sustaining private car use. *Mobilities*, 10(5), 726–747. <https://doi.org/10.1080/17450101.2014.944400>
- Lopez, G. (2018, May 25). ASMR, explained: Why millions of people are watching YouTube videos of someone whispering. *Vox*. <https://www.vox.com/2015/7/15/8965393/asmr-video-youtube-autonomous-sensory-meridian-response>
- Madrigal, A. (2017, August 23). Inside Waymo’s secret world for training self-driving cars. *The Atlantic*. <https://www.theatlantic.com/technology/archive/2017/08/inside-waymos-secret-testing-and-simulation-facilities/537648/>
- Morales, C. (2020, October 7). Millions of views later, Nathan Apodaca keeps the vibe going. *The New York Times*. <https://www.nytimes.com/2020/10/07/us/ocean-spray-buys-car.html>

- Parker, G., Van Alstyne, M., & Jiang, X. (2017). Platform ecosystems: How developers invert the firm. *MIS Quarterly*, 41(1), 255–266. <https://www.jstor.org/stable/26629646>
- Pavlínek, P., & Janák, J. (2007). Regional restructuring of the Škoda auto supplier network in the Czech Republic. *European Urban and Regional Studies*, 14(2), 133–155. <https://doi.org/10.1177/0969776407076101>
- Pine, J., II, & Gilmore, J. H. (2011). *The experience economy*. Harvard Business Review Press.
- Sheller, M. (2004). Automotive emotions - feeling the car. *Theory, Culture & Society*, 21(4–5), 221–242. <https://doi.org/10.1177/0263276404046068>
- Stallman, R. M. (2015). *Free software, free society* (3rd ed.). Free Software Foundation.
- Tesla. (2023). *Autopilot and full self-driving capability features*. Tesla. Retrieved October 1, 2023, from https://www.tesla.com/en_gb/support/autopilot
- Tkacz, N. (2015). *Wikipedia and the politics of openness*. University of Chicago Press.
- Turner, F. (2006). *From counterculture to cyberculture*. University of Chicago Press.
- van der Vlist, F. N. (2022). *The Platform as Ecosystem: Configurations and Dynamics of Governance and Power*. PhD thesis. Utrecht: Utrecht University. <https://doi.org/10.33540/1284>
- van der Vlist, F. N., Helmond, A., Burkhardt, M., & Seitz, T. (2022). API governance: The case of Facebook's evolution. *Social Media + Society*, 8(2), 1–24. <https://doi.org/10.1177/20563051221086228>
- VirtuallyChris. (2020, July 21). The comma.ai community is centered around openpilot and improving the driving experience. *Discord*. <https://discord.com/channels/469524606043160576/539127103853953038/735269223873380442>
- Wark, M. (2004). *The hacker manifesto*. Harvard University Press.
- Wark, M. (2015, May 1). The vectorialist class. *E-flux*. <https://www.e-flux.com/journal/65/336347/the-vectorialist-class/>
- Wark, M. (2019). *Capital is dead, is this something worse?* Verso.
- Zaraczynski, A. (2020, December 23). How lofi hip-hop will inspire new music in 2021. *Forbes*. <https://www.forbes.com/sites/andreazaraczynski/2020/12/23/how-lofi-hip-hop-will-inspire-new-music-in-2021>



8

Resisting Decisions: Coneheads in California

This chapter considers growing resistance to the decision-making capabilities of autonomous vehicles. In this chapter I consider how different groups have responded to the emergence of autonomous vehicles: politically, organisationally, and socially. Up until recently, relatively few concerted, collective efforts to resist the automation of vehicles have occurred. The question is why?

Much of the opposition to date has been limited to ethical or moral debates about sensing and the (mis)recognition of other road users or the (largely erroneous) utilitarian principles considered to underpin autonomous vehicle decision-making (Ganesh, 2017). Artistic and cybersecurity interventions have likewise centred on the disruption of their machine vision capabilities, posing equally hypothetical but decidedly problematic, perceptual challenges for autonomous vehicles and associated systems.

These examples notwithstanding, there has been scant public opposition to autonomous vehicles. Indeed, that whilst similar efforts have been witnessed in other fields where, for example, sensing and (mis)recognition, or utilitarian decisions have yielded unfair or unjust outcomes—say, in respect to algorithmic border control or housing benefit decisions—these have not crossed over into the automotive domain. One question

this chapter tentatively probes is why a ‘counter-mapping’ of the spaces of autonomous driving, as Luis Alvarez León (2019) calls for, has not happened *en masse* to date.

The main body of the chapter, however, focuses on one particular battle around autonomous vehicles that has been waging since 2022, in the city of San Francisco, US. Long a ‘public laboratory’ for tech experimentation, two robotaxi firms, or autonomous vehicle passenger service (AVPS) operators, applied for permission to run limitless services: Cruise and Google/Alphabet’s autonomous vehicle division, Waymo.¹ This chapter chronicles the story of application, opposition, contentious approval, and subsequent suspension of operations of the former, Cruise. In particular, it narrates the involvement of various municipal bodies responsible both for approving permission and opposing operations: the California Department of Motor Vehicles (DMV), the California Public Utilities Commission (CPUC), and the San Francisco County Transportation Authority (SFCTA).

The second part of the chapter introduces a new element into the picture: an anti-AVPS protest group by the name of Safe Street Rebel (SSR). Formed during the pandemic, to campaign for ‘car-free spaces, transit equity, and the end of car dominance’ (Safe Street Rebel, 2023a, p. n.p.) in the San Francisco area, SSR launched a specific anti-autonomous vehicle action in the midst of the regulatory fight. Targeting Cruise and Waymo vehicle’s machine vision capabilities, citizens were encouraged to make use of the most humble of road phenomena: the traffic cone. Placed gently on the hood of a stationary autonomous vehicle—usually in the dead of night—the traffic cone would render the vehicles paralyzed, unable to proceed until the object had been removed. Dubbed the ‘Week of Cone’, the action constituted a landmark moment in the recent history of AVPS: the first on-street, public form of playful resistance to their existence. Both aspects—*regulatory resistance* and *playful resistance*—typify new moments in the development of autonomous vehicles, rendering the future of AVPS—and autonomous vehicles, more generally—less certain than ever.

¹ Throughout this chapter I will generally refer to these as ‘operators’.

Putting Bodies on the Line?

In attempting to answer the question ‘why have we witnessed few efforts to counter the automation of vehicles?’, I want to return to 2013. As media theorist Douglas Rushkoff writes, ‘one December morning in 2013, residents of San Francisco’s Mission District laid their bodies in front of a vehicle to prevent its passage’ (Rushkoff, 2016, p. 1). The vehicle in question had a specific purpose, and ‘although acts of public protest are not unusual in California, this one had an unlikely target: the Google buses used to ferry employees from their homes in the city to the company’s campus in Mountain View, thirty miles away’ (Rushkoff, 2016, p. 1).

As Rushkoff continues he establishes a greater connection between those protesting the buses shuttling employees to Google’s HQ and those sitting on the bus itself. Local residents were angered by the effect Google was having on their neighbourhoods, namely in respect to sky-rocketing housing costs, and employees were burnt out from the demands their employers were making of them. For local residents and activists, laying down in front of the Google shuttle buses, as well as latterly throwing rocks at them in Oakland, was determined to be an easy, and effective, way of resisting Google’s growing power.

Rushkoff understands this act of resistance not as a battle ‘between San Francisco residents and Google employees’ (Rushkoff, 2016, p. 3) nor between ‘the 99 percent and the 1 percent’ (Rushkoff, 2016, p. 3) as such conflict was typically styled following Occupy Wall Street in 2011. Instead, Rushkoff understands such as a battle between ‘everyone’ and a ‘program that promotes *growth* above all else’ (Rushkoff, 2016, p. 3, authors’ emphasis).

Whilst the term ‘platform capitalism’ isn’t used by Rushkoff (2016), Nick Srnicek’s book on the phenomenon was released in the same year (Srnicek, 2016), and much of Rushkoff’s analysis is based on a critique of platforms. As he contends:

We optimized our platforms not for people or even value but growth. So instead of getting more free time, we ended up getting less. Instead of getting more varieties of human expression and interaction, we pushed for

more market-friendly predictability and automation. Technologies were prized most for their ability to extract value from people in terms of ‘eyeball hours’ and the data that could be derived from them. As a result, we have ended up in an always-on digital landscape, constantly pinged by updates and enduring a state of perpetual emergency interruption – what I call ‘present shock’ – previously known only to 911 operators and air traffic controllers. (Rushkoff, 2016, p. 6)

Thus, the protests against the Google buses can be understood as a protest against platform capitalism’s business model, principally based, at the time, on data extraction. What’s interesting is that this analysis still largely holds, even if the contours of the debate have shifted somewhat, more expressly, onto the role of machine learning (ML) and digital labour supporting ‘automation’.

Many of the aspects of the protests against the Google buses translate into plausible opposition to contemporary autonomous vehicle tests, many of which have taken place on public roads, in various forms, over the last seven years. Indeed, Cruise’s faltering robotaxi tests have been taking place in the exact same city, San Francisco, that saw the original Google bus protests in 2013, since 2015 (Marshall, 2022). Yet, until recently, there have been few protests, no laying of bodies down in front of them, and certainly no rock-throwing. Whilst this could be seen as an evolution of protest tactics, coming to the conclusion that Rushkoff does that maybe the residents of San Francisco’s Mission District have more in common with Google employees than they realized, it is also plausibly an acknowledgement of how testing in the wild by big tech firms is unnervingly common. So familiar that any form of public resistance might appear futile.

The lack of public resistance to autonomous vehicle tests is all the more intriguing for the fact that there have been artistic and cybersecurity-based efforts to disrupt the machine vision of autonomous driving in recent years. James Bridle’s ‘Autonomous Trap 001’ installation (2017) is one such example, where the artist placed a nominally ‘autonomous’ car within a rudimentary ‘trap’ consisting of two sets of white chalk markings: one inner solid line and one outer dashed line. As Bridle explains:

What you're looking at is a salt circle, a traditional form of protection – from within or without – in magical practice. In this case it's being used to arrest an autonomous vehicle – a self-driving car, which relies on machine vision and processing to guide it. By quickly deploying the expected form of road markings – in this case, a No Entry glyph, we can confuse the car's vision system into believing it's surrounded by no entry points, and entrap it. (Mufson, 2017, p. n.p.)

Drawing on the more mystical, or spiritual, elements of movement and vision, Bridle's project hinted at the possible ways in which autonomous vehicles might be challenged or opposed. Ben Nassi, an Israeli researcher and former Google employee, similarly developed a range of projects investigating plausible cybersecurity threats to autonomous vehicles. These have included the prospect of so-called phantom attacks using roadside digital billboards, and 'camera spoofing attacks' using projectors to trick driver-assist devices into responding to speed limit signs (Nassi et al., 2020, 2021).

As Nassi et al. (2020) explain, phantom attacks can be considered a 'perceptual challenge' (Nassi et al., 2020, p. 3) for two reasons. Firstly, in that they exploit a so-called validity gap (Nassi et al., 2020, p. 3) that leaves the autonomous vehicle unable to independently verify what is being sensed. If one aspect of an integrated sensing system detects an object, regardless of what other parts of the system may say, the vehicle is minded to take a 'better safe than sorry' approach (Nassi et al., 2020, p. 3) and treat the object as if it was real. Secondly, in that they take advantage of the way ML models detect objects. In this, as Bunz (2019) explains, image-based ML models do not derive meaning (i.e. recognize objects) through understanding the content of images (as humans might) but through a calculation of image properties, such as edges and textures. As Nassi et al. (2020, p. 3) state, 'most object detection algorithms are essentially feature matchers, meaning that they classify objects with high confidence if parts of the object (e.g. geometry, edges, textures) are similar to ... training examples'. If phantom objects match the same geometric, edge-based, and textural properties the image-based models have been trained on, the chances are the system will treat them as real objects.

In the camera spoofing attacks, Nassi et al. (2021) use a similar approach to the above, this time limiting their efforts to using a projector to spoof road sign images. Here, whilst they suggest that attacks of a similar kind can be performed using ‘adversarial’ methods designed to ‘trick ... deep learning classifiers’ (Nassi et al., 2021, p. 1), their approach is both simpler and more effective. Rather than generating entire phantom objects, in this instance Nassi et al. (2021) specifically project different kinds of road signs to be picked up by a specific advanced driver-assistance system (ADAS), the Mobileye 630 PRO. Through various tests, they determine that the system was insensitive to the colour of the road signs, but both the size and shape of the fake signs didn’t attract the attention of the ADAS. Regarding the most critical aspect of their experiment, the speed represented, Nassi et al. (2021) concluded the system responded to an array of numbers, including ‘speed values ... not used in the real world’ (Nassi et al., 2021, p. 3). In short, that if cyber-attackers were to attack a Mobileye 630 PRO equipped vehicle, they could do so by carefully mimicking the style of real road signs and adding any speed value they wanted.

On a different level, Simon Weckert’s ‘Google Maps Hack’ project in 2020 used ‘99 second hand smartphones ... transported in a handcart to generate virtual traffic jams in Google Maps’ (Weckert, 2020). Social navigation app Waze has long been responsible for increasing traffic through residential neighbourhoods, thanks to its algorithmic route-calculation (Littman, 2019), and hackers have previously caused traffic jams in Moscow by mass-ordering taxis to the same location, using the Yandex platform in 2022 (Roth, 2022). All these examples use a variety of mobile devices, mobile data, mobile apps, mobile navigation, and mobile ride-hailing services, to disrupt the otherwise ordinary flow of vehicles through urban environments. Thus, rather than needing to physically lay down in front of the vehicles, each of these cases shows the ease at which a digital form of disruption can affect the movement of vehicles, whether directly (as in the case of Yandex) or indirectly (as in the case of Weckert and Waze). In effect, through the ‘spoofing’ of actual requests and the appearance of user ‘demand’. That some public autonomous vehicle tests are tests of autonomous ride-hailing services suggests that there is at least some scope for similar such actions.

One obvious reason for the relative lack of resistance to autonomous vehicles is that they are difficult to spot in the wild, geographically limited in scope to particular countries (the US) or particular cities/regions (San Francisco, Arizona). Yet, I would argue that these are reasons specifically *for* concerted, collective action, with the ability to likewise build up capacity and ‘test out’ the feasibility of disruptive action. Autonomous vehicle tests, therefore, are similarly tests of/for disruptive action against such tests, establishing the viability, effectiveness, and support for such actions before being able to scale up and out, much like the operators of the vehicle tests themselves are undoubtedly evaluating: ‘first San Francisco, then the world’ being the motto for all Silicon Valley enterprises before them.

These three examples suggest a possible trajectory for anti-autonomous vehicle protests that centre either on disrupting their machine vision capabilities (or lack thereof) or the services upon which they rely. In the rest of this chapter, I will consider the most significant anti-autonomous vehicle event to date that combines both aspects: opposition to the roll-out of AVPS in San Francisco. Focusing on Cruise, the next section considers the legislative battles around the approval of AVPS permits in the city. Strongly opposed by municipal bodies such as the SFCTA, Cruise finally received approval to operate restriction-less robotaxi services—until disaster struck. In the midst of this regulatory fight, a new—and decidedly novel—anti-autonomous vehicle activist group sprung up, demonstrating the ease with which such vehicles could be stopped in their tracks, all with the help of the humble traffic cone. The subsequent section therefore discusses the growing cultural resistance to autonomous vehicles, for which the action in San Francisco provides a likely, playful blueprint.

Regulatory Resistance

In December 2022, following a period of restricted testing, the California DMV granted Cruise permission to run unrestricted passenger services of their autonomous vehicles within San Francisco (Cano, 2022). Cruise CEO Kyle Vogt regarded it as an important step in expanding their operations to the whole city, 24 hours a day, 7 days a week (Vogt, 2022).

Despite this, the decision was not a formality, with CPUC in charge of making the final resolution, as they were when approving Cruise's initial application for 'Phase 1' of their AVPS deployment programme in June 2022 (CPUC, 2022). Approvals from both the DMV and CPUC are required to operate AVPS in San Francisco (CPUC, 2021), bodies separately responsible for automotive regulation and public utility provision across the state of California.

In making their application to CPUC Cruise were expected to submit a so-called advice letter (Cruise, 2022), an 'informal request to furnish service' (CPUC, 2019, p. n.p.) and further charge passengers, by expanding availability across San Francisco.² In Cruise's case, their application was 'suspended for further staff review' (CPUC, 2023a) until May 15, 2023. According to the relevant legislation, anyone can 'file a protest or respond to an advice letter within 20 days of the date of filing of the advice letter' (CPUC, 2019, p. n.p.). Subsequently, and along with 44 other persons, the SFCTA submitted a response (SFCTA, 2023).

In contrast to organizations supporting the application, such as the San Francisco Chamber of Commerce and a tech industry body called the 'Chamber of Progress', the SFCTA opposed the unrestricted expansion of Cruise's operations in the city. The SFCTA, responsible for transport planning across the whole San Francisco County area, considered the broad expansion of Cruise's operations as 'unreasonable in light of the Cruise AV [autonomous vehicle] performance record' (SFCTA, 2023, p. 3), noting principles or areas Cruise were failing to adhere to, namely regarding *incrementalism*, *transparency*, and *reporting metrics*. These constitute the central features of the SFCTA's opposition to Cruise, largely resting on a stated desire to 'collect new data to support incremental expansion evaluation' (SFCTA, 2023, p. 4).

² CPUC operate a tiered system, in which service providers either request amendments to existing permits that do not affect rate changes (Tier 1), affect rate changes or introduce new services (Tier 2), or constitute substantial changes requiring approval by the commission (Tier 3) (CPUC, 2019). Cruise's initial application, as per CPUC guidance on autonomous vehicle deployment programme approvals (CPUC, 2021), required the submission of a Tier 3 advice letter. Cruise's subsequent amendment, thus, only mandated a Tier 2 advice letter be submitted, as they were 'only' requesting an expansion to already-permitted AVPS. Interestingly, during the initial application process, Cruise had unsuccessfully lobbied the CPUC to remove the requirement for autonomous vehicle operators to submit a Tier 2 advice letter if they desired to alter the geographical extent of their operations (CPUC, 2022).

Incrementalism

SFCTA's main concern in the advice letter is that the approval process should adhere to the principle of 'incrementalism'. In other words, that Cruise (and any others) should pass operational performance milestones and thresholds before being allowed to expand their operations. As they note in the letter, 'a series of limited deployments with incremental expansions – rather than unlimited authorizations – offer the best path toward public confidence in driving automation and industry success in San Francisco and beyond' (SFCTA, 2023, p. 4).

Notably, the SFCTA is happy enough for Cruise and other operators to test their vehicles. What the SFCTA requested instead is that operators abide by an overarching developmental principle (incrementalism) seemingly at odds with an approach favoured by Cruise and others. Most notably, the disruptive, anti-regulatory, 'move fast and break things' ethos that resulted in the cessation of Uber ATG's testing programme in Tempe, Arizona, and other locations, after the death of Elaine Herzberg (Hind, 2022a).

As a key constituent of this incrementalist approach, the SFCTA requested that 'new driverless readiness data collection should be required' (SFCTA, 2023, p. 4), offering the SFCTA and relevant parties the ability to scrutinize operations at each stage. The SFCTA deemed this data critical to be able to support an evaluation of operations under their jurisdiction, noting that CPUC currently mandate operators in San Francisco to report nine kinds of operation data on a quarterly basis (CPUC, 2023b). These include total miles travelled by each vehicle per passenger service, vehicle trip miles, vehicle wait time, vehicle occupancy, and four requirements related to Wheelchair Accessible Vehicle (WAV) requests (CPUC, 2023b). Both Cruise and Waymo 'claimed confidentiality for certain portions of its reports' (CPUC, 2023b, p. n.p.) in each of the three quarterly reports for the city's deployment programme, since it began in February 2022.³

Whilst the SFCTA acknowledged that so-called deployment data is 'broader in scope [than pilot programme data], contains detailed trip

³ These were March to May, June to August, and September to November.

level information, and includes information on vehicle charging and public safety incidents’ (SFCTA, 2023, p. 17), CPUC receives no data at all on a critical feature: planned or unplanned stops that obstruct road lanes. Unplanned stops are especially important, as they constitute the kinds of incidents that prevent evidence of obvious performance limitations. These have also drawn the most public scrutiny and have been subject to multiple press reports (Marshall, 2022; Vincent, 2022), questioning the extent to which Cruise vehicles specifically are ready for expansion, what the SFCTA call ‘unlimited authorization’. Documentation, and scrutiny, of unplanned stops is integral to this stated incremental approach—without which neither CPUC nor the SFCTA have any sense of how common unplanned (or indeed planned) stops are. In the SFCTA’s own words, they say that:

To assess how unexpected and unplanned stops obstructing travel lanes impact the transportation network it is critical to know the location and duration of each unplanned stop. San Francisco also recommends using a metric that assesses the rate at which these unplanned stops occur. Given the importance of transit in meeting state climate and equity goals, special consideration should be given to obstructions impacting transit operations. (SFCTA, 2023, p. 18)

The SFCTA provide details of what these unexpected and unplanned stops by Cruise vehicles have looked like across the city, ranging from claims of ‘erratic driving’ noted by callers to 911 and the making of evasive manoeuvres by other drivers around blockages ‘caused by a disabled AV’ (SFCTA, 2023, p. 7). As further noted, these incidents ranged in length from minutes to hours, with 15% of reported cases involving multiple Cruise vehicles clustered in specific locations, providing a veritable assault course for other drivers, public transport providers, and emergency responders (SFCTA, 2023).

In light of these claims, how is the SFCTA confident in making their demand for an incremental developmental approach? The answer is that, following Alvarez León (2019), they committed to provisionally ‘counter-map’ the spaces of autonomous driving by collating and mapping unplanned stop incidents across San Francisco over a seven-month period

(SFCTA, 2023). This provisional counter-mapping—executed in a rather ad-hoc manner—offers an insight into a form of regulatory radicalism offered by a public body such as the SFCTA, willing to provide robust evidence of the present unsuitability of autonomous vehicles for unlimited authorized passenger services on public streets.

What is fascinating is that a public transportation body such as SFCTA has engaged in a counter-mapping process typically executed by non-state actors (Peluso, 1995). In the absence of non-state actors wishing to organize against the disruptive tendencies of operators such as Cruise, a public transportation authority used its comparative power to do the same. In a more critical vein, the SFCTA has not sought to oppose Cruise from using San Francisco as a testbed *tout court*, only to demand operators follow an incremental, evidence-based approach, that does not seek to subject citizens to an increased risk of danger from operations, day and night, all across a city of over 800,000 people.

Transparency

The second concern is that operators should commit to greater levels of transparency. According to the SFCTA, both Cruise and Waymo have ‘sought confidential treatment of basic operational data about AV driving’ having ‘submitted reports to [CPUC] in redacted form’ (SFCTA, 2023, p. 3). Publicly available from a designated online repository (CPUC, 2023b), the reports submitted by Cruise make for indicative reading. In a file marked ‘AV Trips Part 0_REDACTED’ corresponding to all passenger trips made by Cruise vehicles from September 1 to November 30, 2022, all 3352 rows (i.e. 3352 trips) show multiple redacted columns. These include columns specifying trip start date and time, the census tract code of the trip requester, the zip code of the trip requester, dates and times of when the trip was accepted by Cruise, three columns corresponding to vehicle miles travelled during and between trips, and related data on passenger drop-off points and times. In all, 17 columns are labelled as redacted. The only recorded fields of value include trip fulfilment data, trip cancellation data, and passenger total data.

As the SFCTA remark, vehicle miles travelled (VMT) data, currently recorded in three forms through the quarterly reporting mechanism, ‘would provide context for rates of unexpected and unplanned AV stops obstructing travel lanes’ (SFCTA, 2023, p. 19). Yet, as they add, ‘Cruise has redacted data showing VMT of its deployment operations and even high level location information’ (SFCTA, 2023, p. 19) such as census tract and zip code level data of vehicle requesters and passenger drop-off points. In other words, Cruise’s redaction of such operational data is, in the SFCTA’s opinion, a barrier to an incremental approach to the deployment of AVPS in San Francisco, preventing CPUC and other interested parties—including members of the public—from scrutinizing the volume of Cruise vehicles that have initiated unexpected or unplanned stops during their permitted operation.

The SFCTA make two demands. Firstly, that operators submit data on a monthly rather than quarterly basis, making it easier for CPUC to make responsive decisions regarding operator readiness. Then secondly, that permit applications should include the submission of the following information to allow for evaluation:

- All driverless vehicle miles travelled (VMT) for each permit;
- Location and duration of unplanned AV stops (including minimal risk condition) (MRC) and vehicle retrieval events (VRE) obstructing travel lanes by vehicle and underlying permit; and
- Passenger pick-up stops by location, distance from curb, and dwell time for all passenger stops (SFCTA, 2023, p. 19)

Whilst, as noted above, VMT data is currently requested, this is only on a quarterly basis, and Cruise regard this as commercially sensitive data, subsequently redacting all VMT data. MRC is a state autonomous vehicles enter when seeking to reduce risk from collision or another incident, without handing control over to a human operator. Unplanned stops initiated by Cruise will have been the result of specific vehicles or indeed clusters of vehicles as per the reports, entering an MRC state, to minimize the risk of contact with other road users.

If such vehicles cannot be remotely re-activated or the incident resolved from a distance, a vehicle retrieval event (VRE) will be initiated, in order

to return the vehicle to its operational base. Specific passenger pick-up locations, distance from curb, and dwell time for passenger stops are all critical datapoints for a public transportation authority like the SFCTA, mandated to provide safe and secure transportation services for the public at large. Requesting such information on a monthly basis should, the SFCTA argue, be the basis of a permit approval process which, at the moment, is not mandated by CPUC.

As the SFCTA plainly state, ‘although the metrics may reflect on permittee performance in ways that applicants find uncomfortable, none of these data fields call for information that can be legitimately described as protected trade secrets’ (SFCTA, 2023, p. 19), a claim currently made by Cruise and Waymo. Indeed, as they assert, CPUC ‘should require applicants to submit this data in public form without opportunity for claims of confidential treatment’ (SFCTA, 2023, p. 19)—a claim that prevents an incremental approach to deployment and increases the risk to fellow road users where in operation.

Luis Alvarez León establishes ‘three separate possible avenues for counter-mapping’ (Alvarez León, 2019, p. 6) the spaces of autonomous driving, concerning legislation, design, and hacking. The SFCTA’s response to Cruise arguably incorporates aspects of all three. Firstly, most obviously, they are engaging in the formal, legal process through which permits are granted, registering an official protest to an expansion of Cruise’s operating licence, through their response to Cruise’s advice letter. Much like the final resolution issued by CPUC in response to Cruise’s initial application (CPUC, 2022), the commission would be expected to evaluate the protests and responses submitted by interested parties in making their final decision. Secondly—in order to combat closed technological architectures and security threats—the SFCTA explicitly call for Cruise to release operational data records relating to their operations. Whilst the SFCTA are not engaged in, nor advocate, for the hacking of Cruise vehicles, their call for greater transparency and openness of a wide gamut of data produced by them opens the possibility for such a project, whether this involves direct hacking or an indirect form of monitoring in order to record critical data on unplanned stops and similar incidents involving Cruise vehicles.

The SFCTA—alongside other transportation agencies like the Los Angeles Department of Transportation (LADOT) who registered similar concerns—are expressly, as well as indirectly, counter-mapping the spaces of autonomous driving. In short, that it can be said to be offering forms of *regulatory resistance* to the unrestricted expansion of AVPS operations. Here, a form of precautionary development in the guise of ‘technological incrementalism’ fundamentally underpins the counter-strategy, in contrast to the insistent ‘move fast and break things’ logic offered by big tech firms generally and Cruise more specifically. That such a strategy is offered by a municipal agency—and rejected by the firms that would be subject to it—rather than an activist group or collective is striking. It suggests that in the face of the power of big tech—especially in the US and within Silicon Valley itself—public transportation authorities with a broad remit to maintain public transport accessibility and reliability offer a form of resistance that mimics (to some extent) how activist groups protest the state itself. Or, at the very least, that they offer a kind of legislative basis on which activist groups—say, those protesting the expansion of AVPS operations—might offer more radical resistance, echoing earlier efforts to resist Google’s growth in the early 2010s.

Reporting Metrics I: Readiness

Interestingly, the SFCTA use the evidence gathered through the counter-mapping exercise to inform a second part of their incremental approach: the use of a ‘readiness metric’ to assess the readiness of operators for permit approval (SFCTA, 2023). In this, the SFCTA make an important distinction between ‘impact metrics’ and ‘readiness metrics’ (SFCTA, 2023, p. 18). The former, as they suggest, offer a way of evaluating the evolving, anticipated impact of operations. Such impact metrics might reasonably involve the collection of data on the ‘occupancy of AVPS trips, deadheading miles, and their related congestion and energy effects’ (SFCTA, 2023, p. 18).⁴ In other words, these metrics allow CPUC to assess the popularity of such services (occupancy), the immediate impact

⁴ Deadheading miles are those clocked up by vehicles between passenger journeys and return trips to a Cruise logistical base for maintenance or repair.

of such services on local transportation (deadheading miles), and the ongoing, cumulative effect of running operations (congestion, energy).

However, the kinds of readiness metrics desired by the SFCTA operate slightly differently, being less concerned with accumulated impact. Here, the SFCTA are interested in the actual operation of the vehicles in question: where they're going and what they're doing at any one time. More pointedly, the SFCTA want to know how often vehicles are engaging in unplanned stops. Firstly, as discussed in the previous section, this would require operators to collect and share such data with CPUC as part of the permit approval process. However, in a 'raw' form such data is of negligible (or at least lower) value to the authorities, likely unable to determine how significant a total number of unplanned stops is per passenger vehicle, or where the unplanned stops took place.

Thus, to make use of such data, the SFCTA propose a series of readiness metrics it recommends CPUC adopt. These include 'unplanned AV stops obstructing travel lanes in relation to driverless vehicle miles traveled', and 'total lane minutes of obstruction from driverless failures obstructing travel lanes in relation to driverless VMT' (SFCTA, 2023, p. 18). In other words, the SFCTA want to require Cruise and Waymo to make it easier to evaluate how often their vehicles are breaking down. Whilst the first metric offers a way to compare how frequently such services are making unplanned stops, the second metric offers a comparative figure for how serious the resulting obstruction is. Both, when used together, would enable CPUC to properly evaluate—following SFCTA's argument—the 'readiness' of such vehicles for ongoing operations in San Francisco. Indeed, following their demand for developmental incrementalism, such readiness metrics would offer the possibility of granting Cruise, Waymo, and any other operators enhanced permit conditions based on a rising threshold of operability. If these firms were able to improve the operational statistics on unplanned stops and/or unplanned duration, they could be granted permission to operate their vehicles in additional geographical areas and in less restrictive temporal windows.

Reporting Metrics II: Minimal Risk Condition

As the SFCTA highlight, unplanned stops—the empirical object at the centre of their opposition—are a more complicated phenomenon than they would at first appear to be. This is because, as they suggest, stops can be ‘unplanned’ for different reasons.

One such reason is MRC, an operational state entered if a vehicle experiences a technical fault or failure. Entering an MRC is designed to ensure the risk to both passengers and other road users is minimized. In this, an MRC is considered an important ‘failure mitigation strategy’ (SAE International, 2021, p. 10) for autonomous vehicle developers. Except, of course, that entering an MRC state can be considered an unplanned stop for the purposes of AVPS operations, preventing them from carrying paying customers. It is likely that a substantial number of the unplanned stops made by Cruise vehicles reported in the SFCTA’s letter were the result of vehicles entering an MRC state after being unable to perform a ‘dynamic driving task’ (DDT) and unable to hand over to an in-vehicle operator (Fig. 8.1).

In addition, the SFCTA also wish to understand a further element of unplanned stops: not only where vehicles have entered an MRC state but also where they have resulted in human operators, or ‘field staff’ (SFCTA, 2023, p. 7) being called out to (re)move the vehicle itself. VREs constitute an escalation of an MRC, where a vehicle cannot move

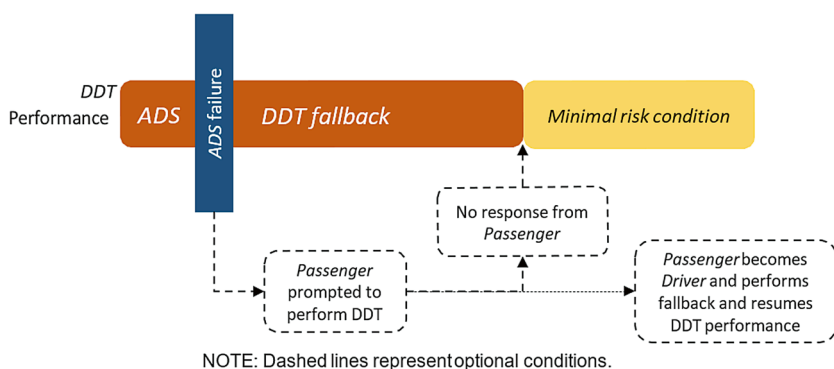


Fig. 8.1 Minimal risk condition. (Source: SAE International, 2021)

autonomously, nor be (re)moved by a remote operator. In such instances, an employee is typically dispatched to retrieve it—necessarily extricating it from the situation and removing it as an obstacle. Naturally, VREs are embarrassing for operators, where retrieval by a human is considered a last resort when both autonomous and (human) remote control have failed. Like before, permittees might find the sharing of such data uncomfortable—although are hardly trade secrets—with events routinely witnessed by fellow road users, uploaded to social media, or reported by local news outlets, as the SFCTA’s letter evidences.

A more straightforward category of unplanned stops also exists: the so-called contact event. Used by Waymo to refer to a range of incidents where contact between an autonomous vehicle and another road user is made, as detailed in Chap. 5, these incidents comprise the bulk of unplanned stops. Although the peculiar ‘mass strandings’ evidenced by the SFCTA are unlikely to be considered contact events, two subsequently reported incidents were: a near-miss collision between one Cruise vehicle and a public bus, after the former had strayed into a bus lane; and an incident where a Cruise vehicle had driven onto a light rail track, blocking a right of way (SFCTA, 2023, pp. 12–15). In both incidents the vehicle had not entered an MRC state, and neither required a VRE. As the SFCTA deduced, both incidents likely constitute violations of the California Vehicle Code (CVC), including stopping in a transit lane (CVC 22500) and not yielding to another vehicle at a right of way (CVC 21800) (SFCTA, 2023, pp. 11, 14).

Evidently, by recommending operators share data on unplanned stops and obstructions across three categories—MRCs, VREs, contact events—as a condition of permit approval, the SFCTA propose a robust, ostensibly workable, route to public AVPS operations. Without such provisions, as the SFCTA patiently argue, greater interruption to the daily lives of San Franciscans is inevitable. Without a commitment to an incremental approach, Cruise and Waymo (quite literally) have a licence to continue operating on whatever terms they wish—much to the wider detriment of public transportation and public life in the city.

Whither Resistance?

Despite the efforts of the SFCTA, on April 25, 2023, CEO Kyle Vogt announced that Cruise had been given permission to operate 24 hours a day, 7 days a week, across the whole of San Francisco (Vogt, 2023a). Although Vogt didn't expressly mention it, the announcement was the result of a draft resolution published by CPUC approving authorization for Cruise's expanded operations (CPUC, 2023c). What Vogt also didn't mention was that the resolution itself was still in draft form, yet to be ratified and adopted by CPUC. Only after this process—subject to a 30-day review period and a 20-day comment period—would Cruise be granted a permit to extend operations.

Vogt clearly regarded the moment as significant: 'operating robotaxis in SF [San Francisco] has become a litmus test for business viability. If it can work here, there's little doubt it can work just about everywhere' (Vogt, 2023b, p. n.p.). Becoming the first company to run relatively unrestricted operations across a major city, with few spatial, temporal, or operational limits, would indeed be momentous. No company in the recent, feverish history of commercial autonomous driving has managed such a feat. Cruise provisionally receiving regulatory approval was something more celebrated rivals, including Waymo and Uber ATG, had not yet managed.

Yet, the significance of Vogt's statement hinged not on Cruise's successful navigation of regulatory requirements. Instead, it was the explicit extrapolation offered by Vogt: 'if it can work here'—in San Francisco—'there's little doubt it can work just about anywhere'. Vogt's logic relies on two implicit truths (a) that San Francisco is the ideal laboratory to test the 'business viability' of AVPS operations and (b) that by extension, that all other cities—whether in the US or elsewhere—are a mere regulatory, social, cultural, and commercial extension of San Francisco. Rather than a city of exceptional exception, San Francisco is an ordinary blueprint for AVPS operations the world over: a 'mundane' model (Hind, 2022b, p. 470) of commercial success.

The draft resolution composed by CPUC is worth considering in detail. It notes that in launching their AVPS 'deployment programme' in

2021, CPUC has endeavoured to deliver four aims: protect passenger safety, expand the benefits of autonomous vehicles to all Californians, improve transportation options for all, and reduce greenhouse gas emissions (CPUC, 2023c). CPUC regards the deployment programme, thus, as a route to achieving at least some aspects of all four of these aims, monitoring permit holders such as Cruise in progressing towards these goals. Central to receiving a permit approval on the deployment programme was a submission by applicants of a 'Passenger Safety Plan' (PSP), submitted by Cruise in December 2022 following their application to expand operations (Cruise, 2022). Comprising of eight general points, any PSP requires applicants to describe how they would minimize safety risks to passengers using their services. After protests and responses were recorded, regarding PSP items, Cruise responded, providing additional measures on 12 specific points of their PSP (Cruise, 2023a), including issues raised by the SFCTA and others in response to the original advice letter submitted by Cruise (2022).

A number of these additional measures focus on expanding the control passengers would have of various functions within the Cruise vehicle during a ride, including 'the ability to have the AV honk and/or extend wait time' (Cruise, 2023a, p. 7). In addition, passengers would be able to review a series of safety messages prior to the trip, displayed both on the in-vehicle touchscreens and the Cruise mobile app (Cruise, 2022, pp. 13–15). Here, much of the adapted PSP focuses both on the 'passenger experience' as well as passenger safety. What these additional measures offer, therefore, is a series of protocols for instilling and governing appropriate user behaviour before, during, and after, a journey in a Cruise vehicle (Fig. 8.2).

Yet, much of the SFCTA's critique of Cruise operations, as detailed before, rest on issues of wider safety of public transportation users not comprehensively covered by Cruise in their response. In this, Cruise simply repeat that their vehicles are 'designed to operate at crossings with cable cars, streetcars, or light rail vehicles' (Cruise, 2022, 24), not that these vehicles have failed to do so on a number of reported occasions. That 'Cruise has designed a thoughtful, integrated system of automated monitoring and response to passenger feedback to appropriately detect and respond to unsafe scenarios outside the vehicle' (Cruise, 2022, p. 46)

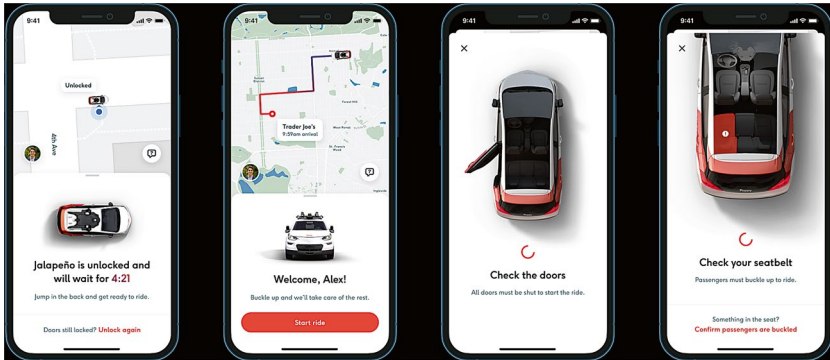


Fig. 8.2 Cruise's 'in-app contextual cues'. (Source: Cruise, 2022)

isn't supported with evidence of such a system, 'thoughtful' or otherwise. In the draft resolution, CPUC simply write that they are 'encouraged by the safety record in passenger services to date' (CPUC, 2023c, p. 11), noting just '5 collisions under its Driverless Deployment permit since receiving its permit in June 2022, none of which resulted in injuries' (CPUC, 2023c, p. 11).

Aside from passenger experience/safety and the wider operation of Cruise services, the draft resolution also refers to the SFCTA's explicit call for an incremental approach. This recommendation formed the foundation of SFCTA's reply to Cruise's original advice letter, in which they note Cruise's non-compliance with permit requirements to make operational data related to trips publicly available.

San Francisco Strikes Back

Barely five weeks later, thanks to another draft resolution issued by CPUC, Waymo were also on the cusp of joining Cruise. This time, the SFCTA joined forces with two other San Francisco bodies: the San Francisco Municipal Transportation Agency (SFMTA) and the San Francisco Planning Department, publishing an official comment on the Waymo draft resolution (San Francisco, 2023). Here, in a repeat of the argument made by the SFCTA in the Cruise case, the trio of San Francisco

bodies stated that ‘it is unreasonable for the Commission [CPUC] to approve Cruise and/or Waymo Advice Letters before adopting expanded reporting requirements and minimum performance standards’ (San Francisco, 2023, p. 5). Whilst the commissioner in charge of the case had recently announced new reporting requirements, the San Francisco bodies argued that approving operations in advance of receiving performance data was ‘inconsistent with the Commission’s power and duty to protect not only passenger safety but the safety of the general public’ (San Francisco, 2023, p. 5).

Here, as with Cruise, the bodies reiterated SFCTA’s incrementalist position, buttressed by evidence of a ‘significant increase in both Waymo Reported Incidents and the cumulative effect of Cruise and Waymo Reported Incidents’ (San Francisco, 2023, p. 5). Uniting behind this common position, the San Francisco bodies argued CPUC had got things the wrong way round. Why allow Cruise and Waymo to operate *tout court*, when the mounting evidence suggests a litany of related hazards and concerns?

Along with the united front against the operators, the trio directly targeted CPUC itself. Despite attempts by the latter to deflect responsibility onto the California DMV—the body responsible for issuing driving licences in the state—the three argue that CPUC ‘has both jurisdiction and a duty to address the hazards raised’ (San Francisco, 2023, p. 8). Indeed, that CPUC ‘should not rely on DMV acquiescence as a basis for [their own] inaction’ (San Francisco, 2023, p. 8). Put succinctly, whilst the DMV’s approval of Waymo’s operational design domain (ODD) ‘sets a ceiling on potential Waymo driverless commercial deployment; it does not set a floor’ (San Francisco, 2023, p. 8). In other words, in the eyes of the trio, it was up to CPUC to establish a foundation for AVPS operators in line with their ‘broad mandate to protect public safety’ (San Francisco, 2023, p. 8). Leaving such decisions in the hands of another body would be an abdication of responsibility.

Beyond the incidents logged by the SFCTA and others, a series of more absurd incidents were noted by the trio. Such incidents included ‘intrusions into marked construction zones in which City employees are working in and under city streets’ (San Francisco, 2023, p. 11) where, for example, on January 13, 2023, ‘a Waymo driverless car drove into the

middle of a construction site and stopped right before rolling into an open trench where San Francisco city employees were working’ (San Francisco, 2023, p. 11). Further incursions ‘into crime scenes and scenes with downed power lines and other hazards marked with caution tape’ (San Francisco, 2023, p. 11) provide evidence that Waymo vehicles are less able to comprehend temporary changes to the driving environment whether they be roadworks, crime scenes, or other unplanned events.

Rather than improving, the San Francisco bodies argued the situation was getting worse. Whilst January 2023 saw five reported incidents involving Waymo vehicles (19 Cruise incidents), February saw a further 10 (19 Cruise). In March 2023, 34 incidents involving Waymo vehicles were reported (59 Cruise), before another 30 in April (57 Cruise) (San Francisco, 2023, p. 14). Aside from any analysis of operator performance, the figures represented something else altogether different. Faced with growing awareness of how Cruise and Waymo were running such services in San Francisco, there appeared to be an emergent resistance movement, ultimately constituting—in collective on-the-ground reporting and engagement with the permit approval process—a form of opposition to the disruptive intentions of AVPS operators. Here, an emergent counter-cartography of autonomous driving was finally taking shape through disputes over the public release (and redaction) of operational data, official responses by affected bodies with a public mandate, and the forceful defence of technological incrementalism.

Playful Resistance: Coneheads II

I began this chapter by wondering why protests against autonomous vehicles hadn’t yet occurred—not least in San Francisco, a city that famously protested shuttle buses ferrying tech workers to the campuses of Google and others, as Douglas Rushkoff (2016) considered. In July 2023 that all changed. Enter an activist group called Safe Street Rebel (SSR) aiming to ‘end car dominance and save the planet one direct action at a time’ (Safe Street Rebel, 2023b, p. n.p.). Whilst not new—they’d formed in opposition to the re-opening of a stretch of San Francisco’s Great Highway in 2021 (Truong, 2023)—they decided to launch of campaign

against Cruise and Waymo, in anticipation of the delayed CPUC decision to expand operations across San Francisco on July 13 (Truong & Mojada, 2023). The group decided to call their campaign the ‘Week of Cone’, in which anonymous members of the group, upon spotting Cruise vehicles conducting trips across the city, were placing—and willing others to place—ubiquitous orange traffic cones on the hoods of the vehicles (Safe Street Rebel, 2023c). Upon doing so, the vehicles appeared bamboozled, hazard lights flashing, seemingly ‘paralyzed’ (Zipper, 2023, p. n.p.), unable to move without a Cruise engineer manually resetting the vehicle system and removing the offending traffic cone. In the fight against AVPS operators, San Francisco seemingly had a new weapon, deployable not in fraught CPUC permit hearings but on the streets, ordinarily under the cover of darkness, when activists were less likely to be seen (or caught) interfering.

Such activity can be seen as part of a broader global movement against car dependency in recent years, from activists like Extinction Rebellion and Just Stop Oil in the UK and *Letzte Generation* in Germany, all of whom have routinely staged road blockades in the name of the climate crisis (Zelden-O’Neill, 2022; Booth, 2023; Stole, 2023), to the Tyre Extinguishers, another activist group working under the cover of darkness to deflate the tyres of SUVs in cities across the world (Gayle, 2022). SSR’s incorporation of anti-AVPS protests in San Francisco into general opposition to car dependency and the climate crisis was arguably novel. As such, it represented a watershed moment long in the making, both in relation to the wider vehicle-facilitated climate crisis, but also connected back to the tech bus protests in San Francisco in 2013. In short, it’s a moment in which a more radical counter-mapping of the spaces of autonomous driving has finally occurred (Alvarez León, 2019).

It is worth expanding a little on what the action entails. As the composite image shows (Fig. 8.3), SSR posted an explainer video on TikTok on how to ‘cone’ an autonomous vehicle (Safe Street Rebel, 2023b). The video itself, just over a minute long, comprises three acts: identifying the problem of AVPS operators in San Francisco, demonstrating the action intended to stop them, and specifying the aim to permanently ban (or at least limit) operators in the city. The video considers the argument by Cruise and Waymo that ‘their cars will reduce traffic and collisions’ before

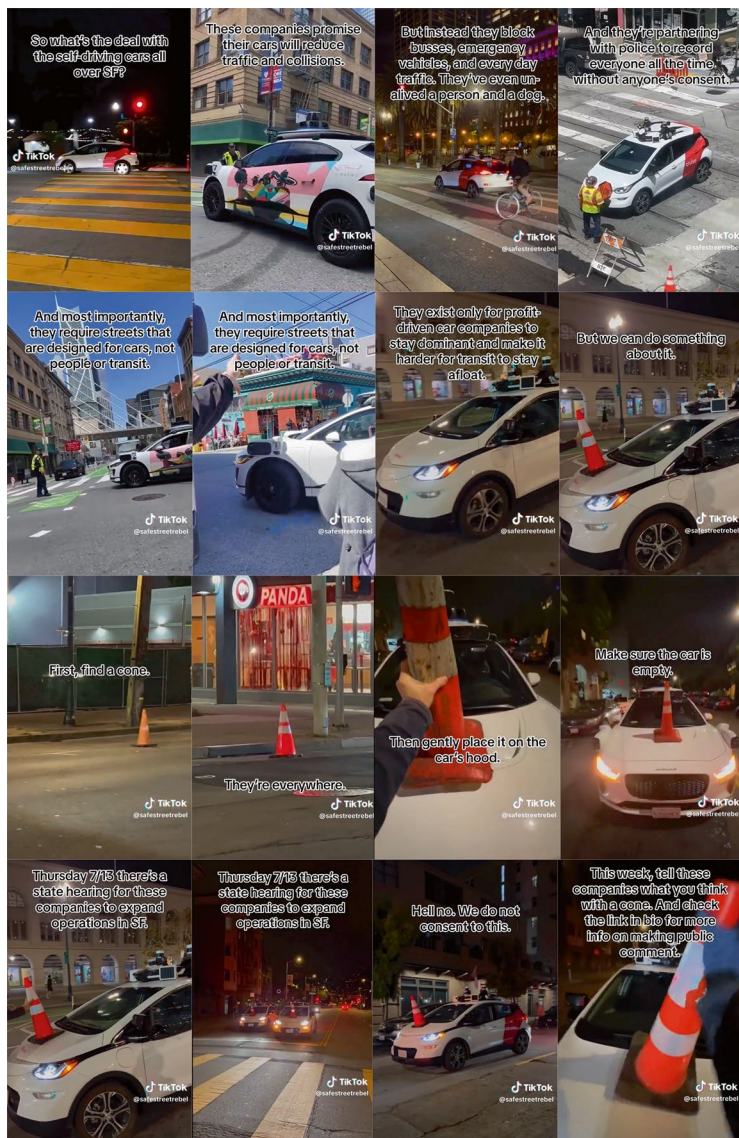


Fig. 8.3 Frames from Safe Street Rebel's TikTok video of how to 'cone' an autonomous vehicle. (Source: Safe Street Rebel, [2023c](#))

arguing they actually ‘block busses, emergency vehicles, and everyday traffic’, an issue widely reported (Marshall, 2022; Vincent, 2022), whilst ‘partnering with police to record everyone all the time without anyone’s consent’. The latter, a story confirmed in a leaked San Francisco Police Department (SFPD) training document (Gordon, 2022), is similar to the backdoor access provided to police departments by Ring doorbell video devices (Belanger, 2022), echoing broader concerns about the ‘functional sovereignty’ (Pasquale, 2017) of autonomous vehicles (Gekker & Hind, 2019). ‘Most importantly’, the video continues, autonomous vehicles ‘require streets that are designed for cars, not people or transit’, echoing SSR’s general opposition to car dependency in San Francisco and across the US.

Then the video turns to the main event: how to ‘do something about it’. First, it instructs viewers to find a humble, ubiquitous, traffic cone (‘they’re everywhere’). Then, upon locating an autonomous vehicle, ‘gently place it on the car’s hood’—making sure ‘the car is empty’ before doing so. Upon placing the cone—as seen in the video—both Cruise and Waymo vehicles automatically come to a halt, activating their orange hazard lights in the process. The placing of the cone by a disembodied arm emerging first from the left of the frame, then POV-style from the bottom, adds a decidable air of humour to the sequence (Fig. 8.4). In the final act viewers are informed of the upcoming CPUC hearing on July



Fig. 8.4 ‘Find a cone (they’re everywhere), then gently place it on the car’s hood’. (Source: Safe Street Rebel, 2023c)

13, in which they are encouraged to offer public comment.⁵ Under the text outlining the details, the video cycles through a series of instances of Cruise vehicles (and in one case, two at the same time) being brought to a halt with a traffic cone. The video ends with a final message: ‘this week, tell these companies what you think with a cone’.

The action brings to mind the playful forms of direct action witnessed across the world in the modern era, from the Situationists in the 1960s, to the anti-capitalist group the ‘Wombles’ in the early 2000s, and Eclectic Electric Collective/Tools for Action in the 2010s. Playful forms of protest have always involved the use—and mis-use—of everyday objects, from maps to inflatables (Hind, 2015), routinely deployed to disrupt everyday actions and movements (Hind, 2016). These so-called disobedient objects (Flood & Grindon, 2014) form the repertoire, and the arsenal, of activist groups the world over, deployable against various enemies from police forces and right-wing groups, to—now of 2023—autonomous vehicles. That an anonymous activist group is issuing a call to social media users to participate in, and proliferate the action, is only novel in respect to the primary platform used (TikTok), as different forms of social media from Indymedia to Blackberry Messenger to the ‘anti-kettling app’ Sukey (Kingsley, 2011; Hind, 2015, 2016) have long been used to disseminate protest-related information and calls to action.

The traffic cone has always been the strange, inanimate foe of the autonomous vehicle. Following Alexis Madrigal’s report on Waymo’s private Californian test facility (Madrigal, 2017), it was questioned whether its vehicles had a trustable value system (Hind, 2019). Noting the ‘prop stash’ Waymo used to test the object-recognition capabilities of their vehicles, I specifically asked ‘what has the right to enter the sacred prop stash? What deserves to be classifiable and therefore recognizable?’ (Hind, 2019, p. 412). From dummies to fake plants, and from skateboards to traffic cones, Waymo was already testing the ability of their vehicles to sense, and react to, any number of miscellaneous items, each with their own specific qualities within a given driving environment. I subsequently considered the likely consequences of *mis-valuing* such objects—not of failing to recognize them wholesale, but wrongly categorizing them

⁵ The hearing was subsequently delayed until August 10.

according to a miscalculation of their assumed form or scrutinized actions. Whilst I suggested that the ‘fisherman’s problem’ (Crampton, 2002; Olsson, 2002) posed issues about the ability for such vehicles to understand the ‘swarming social reality ... to be found far beyond the (sub)urban confines of California’ (Hind, 2019, p. 412), it wasn’t immediately clear they would also be barely unable to comprehend the ‘swarming social reality’—traffic cones and all—of an 8 km² downtown area of San Francisco.

What the cone-based actions show is, less the ability of autonomous vehicles to ‘sense’ or ‘comprehend’ such a social reality and more its inability to compete in it. As the death of Elaine Herzberg also showed, autonomous vehicles might well have the capability to make sense of neatly defined and demarcated objects acting as they should, but far less of an ability to comprehend, and react to, other things *acting together*. For Herzberg, her ‘issue’ (simply her reality) was choosing to walk her bicycle rather than ride it, meaning both human (with a ‘strange’ object beside her) and object (with person interacting with it ‘abnormally’) had had their intrinsic qualities modified somewhat, making something in excess of themselves, as everything tends to do when acting, and inter-acting, in the social world at large. That the humble traffic cone might find itself being re-appropriated, taken out of context, gifted agency, finally at long last animated, was patently too much for an autonomous vehicle and, indeed, its makers, to handle. With no ability to fight back the humble traffic cone and its instigators won—at least, that is, for a fleeting moment.

Decision: Approved

After a six-hour hearing, despite the best efforts of activists and opponents, CPUC voted to approve Waymo and Cruise’s requests to expand operations across the whole of San Francisco, without restrictions (CPUC, 2023d, 2023e, 2023f). It subsequently became the first time, anywhere in the world, that AVPS operators have been permitted free from geographical or temporal conditions. San Franciscans tired of being ‘guinea pigs’ in a city frequently used as a ‘tech-bro playground’ as residents have contended (Bindman, 2023, p. n.p.), now had to prepare

themselves for being round-the-clock ‘test subjects’ (Bindman, 2023, p. n.p.). That CPUC commissioner John Reynolds led the votes in favour of approval did not escape notice by observers. Reynolds, as widely reported (Hawkins, 2023a; Thadani & Merrill, 2023), was a former general counsel of Cruise’s parent company, GM, who had ‘recused’ himself from previous votes.

Drawing on Noortje Marres (2020), the public testing of Cruise and Waymo vehicles has forcefully moved from one (limited, experimental) mode of testing into another (unlimited, experimental). Here, there is no pretence or deployment of technological incrementalism nor evidence-based, precautionary expansion, but a concerted effort (by Cruise and Waymo) to force San Francisco, and San Franciscans, to live with autonomous vehicles. Here, Marres’ notion of ‘co-existence’ (Marres, 2020, p. 549)—a commitment and challenge to which the street trials of autonomous vehicles she studies centred on—appears quaint in this context. There is, in short, no ‘learn[ing] to get along’ (Marres, 2020, p. 549). Indeed, forcing residents and fellow road users to ‘get along’ with the vehicles whether they like it or not, becomes a stark and cynically motivated attack *on*—rather than a more benign test *of*—social life (Marres & Stark, 2020), and the social road (Marres, 2020) itself. As Marres and Stark write, a general shift in the ‘sites and the logics of testing’ (Marres & Stark, 2020, p. 433) has been led by the specific software development cultures fostered by, and in, big tech firms, where the ‘social environment is itself the object of testing’ (Marres & Stark, 2020, p. 433). That this approach has been followed by Google/Alphabet-owned Waymo is no surprise, but its emulation by GM-owned Cruise demonstrates the spread of such a culture into new domains such as the automotive industry, typically governed by very different hardware-based, precautionary safety approaches to automation and control. Whilst ‘platform automobility’ (Hind & Gekker, 2024; Hind et al., 2022) is typically understood through a political-economic lens, constituting a rearrangement of production-based processes and operations according to a platform logic, it becomes more obvious here how this platformization constitutes and solidifies a new, forceful, logic of ‘testing’ too—not simply stretching the limits of acceptable test conditions and procedures but actively destroying them.

CPUC Giveth, California DMV Taketh Away

In another twist to the story and after a recent hit-and-run incident implicating a Cruise vehicle (Hawkins, 2023b), on October 24, 2023, the California DMV suspended Cruise's operations (California DMV, 2023). After two and a half months of unrestricted trips throughout San Francisco, with customers able to call on Cruise 24/7, the company's operations became subject to an indefinite suspension. In response, Cruise took the decision to 'proactively pause driverless operations across all of their fleets' (Cruise, 2023b, p. n.p.) not only in San Francisco but also in Austin, Texas, and Phoenix, Arizona.

Cruise's suspension was based on their contravention of four regulatory points: (a) that a manufacturer's vehicles are not safe for public operation, (b) that a manufacturer has misrepresented information pertaining to the safety of their vehicles, (c) through act or omission that the manufacturer's operations post an unreasonable risk to the public, and (d) that the DMV has the right to suspend or revoke licences due to unsafe practices engaged in by manufacturers (California DMV, 2023). As Kirsten Korosec understood, the suspension of Cruise's permit was specifically the result of Cruise failing to submit video footage of the aftermath of the hit-and-run incident (Korosec, 2023), obstructing efforts to investigate it. In short, the California DMV were no longer convinced of Cruise's ability to run safe operations in San Francisco. Ironically, this failure to submit a form of operational data was precisely the concern put forward by the SFCTA, SFMTA, and the San Francisco Planning Department when calling for an incremental approach.

Interestingly, of course, it was not the California DMV who were responsible, in the first instance, for issuing Cruise with their permit for operating without restriction, on October 24, 2023. Rather, it was CPUC. With a degree of overlapping jurisdiction over the operation of such services—CPUC responsible for public service provision (i.e. 'passenger services') and the California DMV for motor vehicle safety (i.e. autonomous vehicles)—both, in principle, have had a stake in how Cruise operates in San Francisco. The hit-and-run incident that implicated a Cruise vehicle was, it seems, the straw that broke the camel's

back—the last in a long line of incidents involving Cruise vehicles since being given permission by CPUC to run 24/7 operations in August. On November 20, 2023, Cruise CEO, Kyle Vogt, resigned—bringing an end to his 10-year mission to automate driving as the head of the company (Korosec & Bellan, 2023).

Conclusion

This chapter has narrated the ongoing story of the rollout of AVPS operations in San Francisco. With many twists and turns, the story is unlikely far from over. Indeed, it may be only just beginning, spreading across the US, far beyond the confines of the famous tech laboratory of San Francisco (Marshall, 2024). For Cruise, at least the Cruise helmed by CEO Kyle Vogt, the dream appears much further away now than at any point in their short history. Their direct rival in San Francisco, Waymo, stumbles on.

Yet, I began by going back to 2013—still in San Francisco—to consider the impact of the first backlash against big tech. Here, throwing rocks at company buses ferrying engineers from the suburbs to out-of-town tech campuses caught on (Rushkoff, 2016), crystallizing opposition to the emergent power of Google and others. In the years since, localized resistance to the global might of tech firms fell away, undoubtedly connected to the revelations from Edward Snowden and others into state forms of digital surveillance—in which telecommunications firms were complicit (Greenwald, 2013). The implications were of global significance—for those living in the suburbs as San Francisco as much as those living anywhere else. Throwing rocks—or anything else—at buses seemed futile but also misplaced.

Over the course of this period, however, as the development and testing of autonomous vehicles has proceeded, different forms of resistance and opposition have bubbled under the surface. Throughout 2023 it bubbled under the surface no longer, first witnessed as a form of persistent *regulatory resistance* offered by the SFCTA and then, for the first time, as an eruption of broader *cultural* and *playful resistance* in the shape of an anti-car activist group called Safe Street Rebel. Crystallizing in the

aftermath of the Covid-19 pandemic, residents of San Francisco once again saw reason to tackle the physical manifestations of big tech power.

What was fascinating is that, this time, it didn't need a rock. Instead, a humble traffic cone—a weirdly troublesome object for the autonomous vehicle—placed gently on the car's hood. Executed under the cover of darkness, over in a moment, each action constituted a spanner in the works, a bug in the machine, a likely ticket for the engineer, and a nuisance for Cruise and Waymo alike.

In between were the kinds of regulatory battles big tech firms are more than willing, and capable, of fighting (Hind & Seitz, 2024). Finding full-bodied opposition in the form of public transport authorities, like the SFCTA and, ultimately, state departments like the California DMV, was matched, to some degree, by the somewhat agreeable nature of the regulator themselves, CPUC. All the usual tricks were pulled out, from regulatory lobbying efforts to the cynical mobilization of minority users, while operators steadfastly refused to comply with data transparency requirements, non-compliance with which was meant, in theory at least, to disbar them from operating. In their hurry to launch, they eschewed an incrementalist approach that might have offered an eventual, sustainable, agreeable route to operation—but big tech firms and those who seek to emulate them, still, despite protestations, like to 'move fast and break things' (Levy, 2014).

Whether this is the shape of things to come is yet to be seen. But what is obvious is that coordinated opposition to eventual rollout of autonomous vehicles—namely in the form of AVPS operations—is here. The counter-mapping of autonomous driving has begun in earnest. How long it needs to continue for, appears—at this moment—to be an open question.

Returning to the opening chapter, I stated that the book was about the phenomenon of autonomous driving—not what is involved from a strictly 'technical' perspective but where the dream of automating driving had led us in recent years. The quest to automate driving has indeed—as the chapters in the book have hopefully shown—been complex, costly, and contentious. From the mapping of ODDs and the simulation of vehicle trajectories to the rise of open-source driver-assistance devices and anti-autonomous vehicle protests, the journey has not been smooth. As

this final chapter suggests, the dream might even have run out of road altogether.

Driving metaphors aside, the most enduring of all relevant metaphors—of the AI ‘pipeline’—is the one I have used to structure the book throughout. I have done so in order to help make sense of the ‘technical work’ being performed at all stages of the autonomous vehicle development process. In so doing, I have sought to focus on the different kinds of decisions necessarily made to grant autonomous vehicles their various perceptive capabilities. Through the lens of the decision, and how decisions are mapped, sensed, trained, simulated, and secured, as well as ‘relaxed’ and resisted, the book has intended to offer an analysis of what Luciana Parisi termed ‘technological decisionism’ (Parisi, 2017). In a world increasingly driven by a logic of technological decisionism, it does not matter that decisions are wrong, only that they are executed quickly, clearly, definitively, and authoritatively. The consequences of this technological decisionism, as I have hopefully demonstrated, are manifold and pervasive.

In the cases I have examined in the book—from the mapping of ODDs in Arizona, to the battles over AVPS permits in California—technological decisionism seeks to override all other value systems and logics. Indeed, as these cases have illustrated, proponents have also sometime sought to actively destroy them. Despite this, state bodies seeking precautionary approaches to the development of AVPS operations, libertarian projects desiring a hacker-ish vision of the automotive future, or activist groups taking to the streets to protest car culture have challenged the power—and vision—of big tech in different ways.

Yet, despite these manifest, and lively, forms of opposition to technological decisionism, the quest to automate driving has faced a more familiar, consistent, foe throughout: the evident technical limits of machine learning and machine vision. As the chapters in this book have hopefully demonstrated, designing vehicles with the capability to see, move, and make decisions autonomous of humans has proved a significant, ongoing challenge. Those challenges have involved tackling a litany of technical, organizational, and logistical hurdles to deliver larger training datasets, more diverse training datasets, better annotated training datasets, more sensors, different assemblages of sensors, more powerful semiconductor

chips, optimized object-recognition techniques, broader simulation parameters, and expanded conflict typologies, amongst many others.

In the process, a great many innovations have emerged, deployed in a real-world domain for the first time. This book has sought to examine these developments—sometimes on their own terms, using the language of those involved, whilst sometimes deploying different analytical frames to make sense of their wider social, cultural, and political impact. Ultimately, however, in machine learning’s ‘quest for agency’, as Roberge and Castelle (2020, p. 13) have put it, such technical innovations have had to—and often failed to—bear responsibility. It is here where the dream of autonomous driving has turned into a nightmare of sorts, as the full pressure of what it means to autonomously make decisions, and to properly make sense of the world, begins to finally weigh heavy.

References

- Alvarez León, L. F. (2019). Counter-mapping the spaces of autonomous driving. *Cartographic Perspectives*, 92, 10–23. <https://doi.org/10.14714/CP92.1450>
- Belanger, A. (2022, July 14). Amazon finally admits giving cops Ring doorbell data without user consent. *Ars Technica*. <https://arstechnica.com/tech-policy/2022/07/amazon-finally-admits-giving-cops-ring-doorbell-data-without-user-consent/>
- Bindman, A. (2023, August 8). San Franciscans are tired of being used as Cruise, Waymo’s ‘test subjects’. *SF Gate*. <https://www.sfgate.com/tech/article/san-franciscans-beg-california-stop-cruise-waymo-18286073.php>
- Booth, R. (2023, January 1). Extinction Rebellion announces move away from disruptive tactics. *The Guardian*. <https://www.theguardian.com/world/2023/jan/01/extinction-rebellion-announces-move-away-from-disruptive-tactics>
- Bridle, J. (2017). Autonomous Trap 001. *James Bridle*. <https://jamesbridle.com/works/autonomous-trap-001>
- Bunz, M. (2019). The calculation of meaning: On the misunderstanding of new artificial intelligence as culture. *Culture, Theory and Critique*, 60(3–4), 1–16. <https://doi.org/10.1080/14735784.2019.1667255>
- California Department of Motor Vehicles. (2023, October 24). DMV statement on Cruise LLC suspension. *California DMV*. <https://www.dmv.ca.gov/portal/news-and-media/dmv-statement-on-cruise-llc-suspension/>

- California Public Utilities Commission (CPUC). (2019). *Revised advice letter process effective January 1, 2020*. CPUC. Retrieved October 1, 2023, from <https://www.cpuc.ca.gov/industries-and-topics/internet-and-phone/advice-letter-information>
- California Public Utilities Commission (CPUC). (2021). *Application guidance for CPUC autonomous vehicle programs*. CPUC. Retrieved October 1, 2023, from <https://www.cpuc.ca.gov/-/media/cpuc-website/divisions/consumer-protection-and-enforcement-division/documents/tlab/av-programs/cpuc-av-program-applications-guidance-20211026.pdf>
- California Public Utilities Commission (CPUC). (2022). *Resolution approving Cruise LLC's application for phase 1 driverless autonomous vehicle passenger service deployment program*. CPUC. Retrieved October 1, 2023, from <https://www.cpuc.ca.gov/-/media/cpuc-website/files/uploadedfiles/cpucwebsite/content/licensing/autovehicle/tl-19137-cruise-driverless-deployment-final-resolution.pdf>
- California Public Utilities Commission (CPUC). (2023a). *Phase I driverless autonomous vehicle deployment program advice letter status*. CPUC. Retrieved October 1, 2023, from <https://www.cpuc.ca.gov/regulatory-services/licensing/transportation-licensing-and-analysis-branch/autonomous-vehicle-programs/phase-i-driverless-autonomous-vehicle-deployment-program-advice-letter-status>
- California Public Utilities Commission (CPUC). (2023b). *AV program quarterly reporting*. CPUC. Retrieved October 1, 2023, from <https://www.cpuc.ca.gov/regulatory-services/licensing/transportation-licensing-and-analysis-branch/autonomous-vehicle-programs/quarterly-reporting>
- California Public Utilities Commission (CPUC). (2023c). *Draft resolution: Resolution approving authorization for Cruise LLC's expanded service in autonomous vehicle passenger service phase I driverless deployment program*. CPUC. Retrieved October 1, 2023, from <https://docs.cpuc.ca.gov/SearchRes.aspx?docformat=ALL&docid=508669474>
- California Public Utilities Commission (CPUC). (2023d, August 10). CPUC approves permits for Cruise and Waymo to charge fares for passenger service in San Francisco. CPUC. <https://www.cpuc.ca.gov/news-and-updates/all-news/cpuc-approves-permits-for-cruise-and-waymo-to-charge-fares-for-passenger-service-in-sf-2023>
- California Public Utilities Commission (CPUC). (2023e). *Resolution approving authorization for Cruise LLC's expanded service in autonomous vehicle passenger service Phase I driverless deployment program*. CPUC. Retrieved October 1, 2023, from <https://docs.cpuc.ca.gov/PublishedDocs/Published/G000/M516/K812/516812218.PDF>

- California Public Utilities Commission (CPUC). (2023f). *Resolution approving Waymo LLC's application for Phase I driverless autonomous vehicle passenger service deployment program*. CPUC. Retrieved October 1, 2023, from <https://docs.cpuc.ca.gov/PublishedDocs/Published/G000/M516/K812/516812344.PDF>
- Cano, R. (2022, December 16). Cruise poised to offer driverless taxi rides to all of San Francisco, 24-7. Here's what to expect. *San Francisco Chronicle*. <https://www.sfchronicle.com/bayarea/article/cruise-poised-to-offer-driverless-taxi-rides-to-17657865.php>
- Crampton, J. (2002). Thinking philosophically in cartography: Toward a critical politics of mapping. *Cartographic Perspectives*, 42(1), 12–31. <https://doi.org/10.14714/CP41.561>
- Cruise. (2022). *Cruise application for driverless deployment permit ODD expansion – Tier 2 advice letter*. CPUC. Retrieved October 1, 2023, from <https://www.cpuc.ca.gov/-/media/cpuc-website/divisions/consumer-protection-and-enforcement-division/documents/tlab/av-programs/2022-12-16-final-tier-2-advice-letter-for-cruise-driverless-deployment-expansion.pdf>
- Cruise. (2023a). *Cruise LLC's reply to protests and comments to Cruise's application for driverless deployment permit ODD expansion*. CPUC. Retrieved October 1, 2023, from <https://www.cpuc.ca.gov/-/media/cpuc-website/files/uploaded-files/cpucwebsite/content/licensing/autovehicle/2023-02-01-cruise-reply-to-protests-and-comments-to-tier-2-advice-letter%2D%2D-0002.pdf>
- Cruise. (2023b, October 27). (1/3) The most important thing for us right now is to take steps to rebuild public trust [...]. *Twitter*. <https://twitter.com/Cruise/status/1717707807460393022>
- Flood, C., & Grindon, G. (2014). *Disobedient objects*. V&A Publishing.
- Ganesh, M. I. (2017). Entanglement: Machine learning and human ethics in driver-less car crashes. *A Peer-Reviewed Journal About Machine Research*, 6(1), 76–87. <https://doi.org/10.7146/aprja.v6i1.116013>
- Gayle, D. (2022, November 29). Tyre Extinguishers deflate tyres of 900 SUVs in 'biggest ever action'. *The Guardian*. <https://www.theguardian.com/environment/2022/nov/29/tyre-extinguishers-deflate-tyres-of-900-suvs-in-biggest-ever-action>
- Gekker, A., & Hind, S. (2019). Infrastructural surveillance. *Big Data & Society*, 22(8), 1414–1436. <https://doi.org/10.1177/1461444819879426>
- Gordon, A. (2022, May 11). San Francisco police are using driverless cars as mobile surveillance cameras. *Vice*. <https://www.vice.com/en/article/v7dw8x/san-francisco-police-are-using-driverless-cars-as-mobile-surveillance-cameras>

- Greenwald, G. (2013, June 6). NSA collecting phone records of millions of Verizon customers daily. *The Guardian*. <https://www.theguardian.com/world/2013/jun/06/nsa-phone-records-verizon-court-order>
- Hawkins, A. J. (2023a, August 11). Robotaxis score a huge victory in California with approval to operate 24/7. *The Verge*. <https://www.theverge.com/2023/8/10/23827790/waymo-cruise-cpuc-vote-robotaxi-san-francisco>
- Hawkins, A. J. (2023b, October 3). Cruise says a hit-and-run ‘launched’ pedestrian in front of one of its robotaxis. *The Verge*. <https://www.theverge.com/2023/10/3/23901233/cruise-crash-hit-run-pedestrian-injury-sf-robotaxi>
- Hind, S. (2015). Maps, kettles and inflatable cobblestones: The art of playful disruption in the city. *Media Fields Journal: Critical Explorations in Media and Space*, 9. <http://mediafieldsjournal.squarespace.com/tactical-frivolity-and-disobed/>
- Hind, S. (2016). *Disruptive cartographies: Manoeuvres, risk and navigation*. PhD thesis. Coventry: University of Warwick. <http://webcat.warwick.ac.uk/record=b3067889~S15>
- Hind, S. (2019). Digital navigation and the driving-machine: Supervision, calculation, optimization, and recognition. *Mobilities*, 14(4), 401–417. <https://doi.org/10.1080/17450101.2019.1569581>
- Hind, S. (2022a). Machinic sensemaking in the streets: More-than-lidar in autonomous vehicles. In G. Rose (Ed.), *Seeing the city digitally: Processing urban space and time* (pp. 57–80). Amsterdam University Press. <https://doi.org/10.2307/j.ctv2j6xrs3.6>
- Hind, S. (2022b). Making decisions: The normal interventions of Nissan ‘mobility managers’. *Mobilities*, 17(4), 467–483. <https://doi.org/10.1080/17450101.2021.1988682>
- Hind, S., & Gekker, A. (2024). Automotive parasitism: Examining Mobileye’s ‘car-agnostic’ platformisation. *New Media & Society*, 26(7), 3707–3727. <https://doi.org/10.1177/14614448221104209>
- Hind, S., Kanderske, M., & van der Vlist, F. (2022). Making the car ‘platform ready’: How big tech is driving the platformization of automobility. *Social Media + Society*, 8(2), 1–13. <https://doi.org/10.1177/20563051221098697>
- Hind, S., & Seitz, T. (2024). Cynical technical practice: From AI to APIs. *Convergence*, 30(1), 29–48. <https://doi.org/10.1177/13548565221133248>
- Kingsley, P. (2011, February 3). Inside the anti-kettling HQ. *The Guardian*. <https://www.theguardian.com/uk/2011/feb/02/inside-anti-kettling-hq>

- Korosec, K. (2023, October 24). California DMV immediately suspends Cruise's robotaxi permit. <https://techcrunch.com/2023/10/24/dmv-immediately-suspends-cruises-robotaxi-permit-in-california>
- Korosec, K., & Bellan, R. (2023, November 20). Cruise co-founder and CEO Kyle Vogt resigns. *TechCrunch*. <https://techcrunch.com/2023/11/19/cruise-co-founder-and-ceo-kyle-vogt-resigns/>
- Levy, S. (2014, April 30). Mark Zuckerberg on Facebook's future, from virtual reality to anonymity. *Wired*. <https://www.wired.com/2014/04/zuckerberg-f8-interview/>
- Littman, J. (2019, August 20). Waze hijacked LA in the name of convenience. Can anyone put the genie back in the bottle? *Los Angeles Magazine*. <https://www.lamag.com/citythinkblog/waze-los-angeles-neighborhoods/>
- Madrigal, A. C. (2017, August 23). Inside Waymo's secret world for training self-driving cars. *The Atlantic*. <https://www.theatlantic.com/technology/archive/2017/08/inside-waymos-secret-testing-and-simulation-facilities/537648/>
- Marres, N. (2020). Co-existence or displacement: Do street trials of intelligent vehicles test society? *The British Journal of Sociology*, 71(3), 537–555. <https://doi.org/10.1111/1468-4446.12730>
- Marres, N., & Stark, D. (2020). Put to the test: For a new sociology of testing. *The British Journal of Sociology*, 71(3), 423–443. <https://doi.org/10.1111/1468-4446.12746>
- Marshall, A. (2022, July 8). Cruise's robot car outages are jamming up San Francisco. *Wired*. <https://www.wired.com/story/cruises-robot-car-outages/>
- Marshall, A. (2024, January 4). Cruise was asked to ground robotaxis on Halloween to keep kids safe. *Wired*. <https://www.wired.com/story/robotaxis-cities-cruise-waymo-officials-regulation>
- Mufson, B. (2017, March 18). Meet the artist using ritual magic to trap self-driving cars. *Vice*. <https://www.vice.com/en/article/qkmeyd/meet-the-artist-using-ritual-magic-to-trap-self-driving-cars>
- Nassi, B., Nassi, D., Ben-Netanel, R., & Elovici, Y. (2021). Spoofing Mobileye 630's video camera using a projector. Preprint. https://www.ndss-symposium.org/wp-content/uploads/autosec2021_23007_paper.pdf
- Nassi, B., Nassi, D., Ben-Netanel, R., Mirksy, Y., Drokina, O., & Elovici, Y. (2020). Phantom of the ADAS: Phantom attacks on driver-assistance systems. Preprint. <https://eprint.iacr.org/2020/085.pdf>

- Olsson, G. (2002). Glimpses. In P. Gould & F. R. Pitts (Eds.), *Geographical voices: Fourteen autobiographical essays* (pp. 237–268). Syracuse University Press.
- Parisi, L. (2017, October 1). Reprogramming decisionism. *E-flux*. <https://www.e-flux.com/journal/85/155472/reprogramming-decisionism/>
- Pasquale, F. (2017, December 6). From territorial to functional sovereignty. *Law and Political Economy*. <https://lpeproject.org/blog/from-territorial-to-functional-sovereignty-the-case-of-amazon/>
- Peluso, N. L. (1995). Whose woods are these? Counter-mapping forest territories in Kalimantan. *Indonesia. Antipode*, 27(4), 383–406. <https://doi.org/10.1111/j.1467-8330.1995.tb00286.x>
- Roberge, J., & Castelle, M. (2020). Toward an end-to-end sociology of 21st-century machine learning. In J. Roberge & M. Castelle (Eds.), *The cultural life of machine learning: An incursion into critical AI studies* (pp. 1–30). Palgrave Macmillan. https://doi.org/10.1007/978-3-030-56286-1_1
- Roth, E. (2022, September 3). Hackers caused a massive traffic jam in Moscow using a ride-hailing app. *The Verge*. <https://www.theverge.com/2022/9/3/23335694/hackers-traffic-jam-russia-moscow-ride-hailing-app-yandex-taxi>
- Rushkoff, D. (2016). *Throwing rocks at the Google Bus: How growth became the enemy of prosperity*. Penguin.
- SAE International. (2021). *Taxonomy and definitions for terms related to driving automation systems for on-road motor vehicles*. SAE International. Retrieved October 1, 2023, from https://doi.org/10.4271/J3016_202104
- Safe Street Rebel. (2023a). *Safe Street Rebel: Not without a fight*. Safe Street Rebel. Retrieved October 1, 2023, from <https://www.safestreetrebel.com/>
- Safe Street Rebel. (2023b). *Safe Street Rebel*. Twitter. Retrieved October 1, 2023, from <https://twitter.com/safestreetrebel>
- Safe Street Rebel. (2023c, July 5). Welcome to week of Cone. *Twitter*. <https://twitter.com/SafeStreetRebel/status/1676695665278337024>
- San Francisco. (2023). *San Francisco comments on the draft resolution approving authorization for Waymo Autonomous Vehicle Passenger Service Phase 1 Driverless Deployment Program*. CPUC. Retrieved October 1, 2023, from <https://www.cpuc.ca.gov/-/media/cpuc-website/divisions/consumer-protection-and-enforcement-division/documents/tlab/av-programs/tl-19144-comments.zip>
- San Francisco County Transportation Authority (SFCTA). (2023). *Re: Protest of Cruise LLC Tier 2 advice letter (0002)*. CPUC. Retrieved October 1, 2023, from <https://www.cpuc.ca.gov/-/media/cpuc-website/files/uploadedfiles/>

cpucwebsite/content/licensing/autovehicle/cruise-al-02-protests-and-responses%2D%2D-revised.zip

Srnicek, N. (2016). *Platform capitalism*. Wiley.

Stole, B. (2023, February 1). Glue-ten Tag! Behind the scenes with Germany's reviled 'climate-gluer' activists. *Slate*. <https://slate.com/news-and-politics/2023/02/letzte-generation-last-generation-germany-climate-gluers.html>

Thadani, T., & Merrill, J. B. (2023, August 10). California just opened the floodgates for self-driving cars. *The Washington Post*. <https://www.washingtonpost.com/technology/2023/08/10/san-francisco-robotaxi-approved-waymo-cruise/>

Truong, K. (2023, July 7). We spoke to one of the activists 'coning' Cruise and Waymo robotaxis in San Francisco. *The San Francisco Standard*. <https://sfstandard.com/2023/07/07/we-spoke-to-one-of-the-activists-coning-cruise-and-waymo-robotaxis-in-san-francisco>

Truong, K., & Mojadad, I. (2023, June 27). Robotaxi service expansion in San Francisco delayed by state regulators. *The San Francisco Standard*. <https://sfstandard.com/2023/06/27/vote-to-expand-cruise-waymo-robotaxis-in-san-francisco-delayed-amid-growing-pushback/>

Vincent, J. (2022, July 1). Driverless Cruise robotaxis stop working simultaneously, blocking San Francisco street. *The Verge*. <https://www.theverge.com/2022/7/1/23191045/cruise-robotaxis-driverless-roadblock-san-francisco>

Vogt, K. (2022, December 16). Exciting news! @Cruise has received approval from the DMV to expand our commercial deployment area to the full 7x7 of San Francisco, 24/7. *Twitter*. <https://twitter.com/kvogt/status/1603548091772465154>

Vogt, K. (2023a, April 25). Well folks, we did it. I have been waiting for this day for almost 10 years. I am proud to announce @Cruise is now running 24/7 across all of San Francisco! This is a pivotal moment for our business. Let me tell you why (1/6). *Twitter*. <https://twitter.com/kvogt/status/1650845475195559939>

Vogt, K. (2023b, April 25). Operating robotaxis in SF [San Francisco] has become a litmus test for business viability. If it can work here, there's little doubt it can work just about everywhere. *Twitter*. <https://twitter.com/kvogt/status/1650845477112516609>

Weckert, S. (2020, December 29). Hacking Google Maps. *Chaos Computer Club*. https://media.ccc.de/v/rc3-197911-hacking_google_maps

- Zelden-O'Neill, S. (2022, October 29). Just Stop Oil activists dragged out of road by motorists in London. *The Guardian*. <https://www.theguardian.com/environment/2022/oct/29/just-stop-oil-activists-dragged-out-of-way-by-motorists-at-london-protests>
- Zipper, D. (2023, August 9). San Francisco has a problem with robotaxis. *The Atlantic*. <https://www.theatlantic.com/ideas/archive/2023/08/robotaxis-san-francisco-self-driving-car/674956/>

Index¹

A

Advanced driver-assistance systems (ADAS), 30, 31, 48, 55, 159, 172, 182, 193, 200–202, 219, 230

Agamben, Giorgio, 160–163, 186

Agre, Philip, 195, 196

Alvarez León, Luis F., 196, 226, 234, 237, 247

Amoore, Louise, 4, 15–18, 99, 126, 133, 142, 150–152, 154, 197

Annotation, 1, 32, 52, 63, 70, 73–75, 77, 82, 85, 88

Argo AI, 2, 33, 113, 115, 159, 205

Autonomous Land Vehicle (ALV), 25, 26, 47, 48

Autonomous vehicle passenger services (AVPS), 31, 35, 36, 226, 231, 232, 232n2, 236, 238, 240–242, 245–247, 251, 254–256

Average precision (mean), 113

B

Beer, Stafford, 6, 7, 129

Benchmark(s), 26, 27, 32, 46, 69, 79–89

phase, 26, 32, 46

Bowker, Geoffrey C., 72, 78

Bratton, Benjamin, 34, 161–164, 167, 174, 186

¹Note: Page numbers followed by 'n' refer to notes.

C

- Calculation, 2, 9, 46, 51, 54, 78, 117, 125, 126, 128, 133, 140, 142, 143, 145, 146, 151–153, 174, 185, 229
- Calibration, 18, 99, 102, 194, 205–209, 213, 216, 219
- Californian ideology, 35, 193–196, 198, 200, 219, 220
- California Public Utilities Commission (CPUC), 226, 232–239, 232n2, 242–245, 247, 249, 251–255
- Classification, 5, 10, 16, 32, 34, 69, 72, 78, 95, 117, 127, 137–139, 148, 152, 154
- Comma, 3, 14, 26, 31, 35, 57, 193, 194, 196–220
- Conflict typology, 34, 146, 147, 257
- Contact events, 33, 34, 125, 137–141, 144–146, 148–153, 241
- Cruise, 1, 27, 35, 36, 105, 146, 159, 163, 167, 211, 212, 226, 228, 231–247, 232n2, 238n4, 249–255

D

- Decisions
 - decisionism, 1–36, 256
 - decision-making, 2–17, 22–24, 31, 33, 34, 44, 57, 62, 96, 101, 108, 126, 129, 132, 134, 137, 140, 146, 152, 154, 160, 162, 164, 166, 187, 195, 225
 - deferred, 11–13, 34, 128, 150–154

theories of, 4, 5, 8–11

- Defense Advanced Research Projects Agency (DARPA), 25, 26, 47, 48, 81, 84, 99, 105
- Domains, 5, 6, 15, 22, 32, 35, 44, 49–54, 56, 62, 67–73, 75, 77, 78, 81–83, 85, 90, 103, 105, 110, 120, 141, 198, 203, 205, 225, 245, 252, 257
- Dynamic scheduling, 33, 114–119

E

- End-to-end approach (E2E), 56, 57, 59, 60, 106

F

- Forecasting, 10, 79, 84, 106, 107, 114, 117, 119, 128

G

- Geometrics, 32, 43, 44, 46, 49, 52, 55–57, 60, 61, 229
- Ground truth
 - cartographic, 48, 75, 76
 - machine learning, 49, 67, 75, 76, 151

H

- Hacker class, 35, 198, 199, 203
- Hong, Sun-ha, 12, 14, 140
- Hotz, George, 193, 197, 200–204, 219

I

Incremental(ism), 3, 26, 27,
 30–33, 35, 46, 69,
 85–89, 232–239, 241,
 244, 246, 252, 253
 phase, 3, 26, 27, 32, 46
 Interesting-ness, 32, 69, 83, 84
 Interoperability, 22, 33, 71, 72, 96,
 98, 100–109, 111–114, 117,
 119, 121
 finessing, 33, 98

J

Jaton, Florian, 11, 75–77, 88

K

KITTI Vision Benchmark Suite, 26,
 32, 68, 69, 79–89, 95

L

Leigh Star, Susan, 3, 72, 78

M

Machine learning (ML), 4–6, 9,
 10, 14–17, 19, 20, 22–24,
 26, 33, 36, 49, 60, 67, 68,
 70–72, 75, 76, 78–80, 83,
 89, 98, 110, 151, 153, 197,
 228, 256, 257
 models, 1, 3, 10, 17, 18, 21,
 32, 44, 49, 50, 52, 57, 59,
 60, 63, 67–69, 71, 76–78,
 88–90, 95, 96, 101, 113,
 209, 229

Machine vision, 2, 4–6, 4n2, 8, 9,
 14, 23, 32, 33, 69, 73, 75,
 78, 79, 82, 90, 95, 96, 100,
 103, 105, 108–111, 118,
 120, 225, 226, 228, 229,
 231, 256

Mackenzie, Adrian, 10, 22, 50,
 72, 97, 101–103, 113,
 116, 119

Mapping, 1, 2, 10, 18, 22, 23, 31,
 32, 43–63, 72, 106, 132, 234,
 255, 256

mapless-ness, 43–63

Marres, Noortje, 21, 22, 140,
 142, 252

Metrics, 75–77, 82, 83, 89, 113,
 232, 234, 237–241

N

National Highway Traffic Safety
 Administration (NHTSA),
 145, 197, 206, 219

National Transportation Safety Board
 (NTSB), 27, 29, 30, 50n1,
 127, 128, 140, 167n2

Navigation, 46–48, 230, 242
 automotive, 43, 45–49

O

Object-recognition, 10, 33, 69,
 77, 79, 89, 90, 96, 100,
 101, 112, 113, 184, 202,
 209, 250, 257

Operations

 operational analysis, 18, 20–23
 operational exceptionality, 164

P

Parisi, Luciana, 14, 15, 97, 98, 256
 Pasquinelli, Matteo, 6, 16, 17, 20,
 21, 49, 60
 Path predictions/trajectories, 57,
 111, 126–128, 133–137, 143,
 148, 154
 Perception, 10, 20, 24, 70, 78, 79,
 83, 85, 95, 96, 106, 107, 109,
 111, 114–117, 119, 120,
 133, 153
 Pipeline, 10, 14, 23–24, 32, 33, 60,
 71, 72, 79, 101, 113, 114,
 137, 198, 256
 Planning, 10, 22, 56, 72, 78, 79, 85,
 89, 90, 106, 107, 114, 117,
 165, 184, 232

R

Robotic phase, 25
 Rushkoff, Douglas, 227, 228,
 246, 254

S

Safe Street Rebel (SSR), 3, 35, 226,
 246–249, 254
 San Francisco County Transportation
 Authority (SFCTA), 226,
 231–245, 253–255
 Semantics, 43, 44, 46, 49, 52,
 54–57, 60, 61, 71, 79
 Semiconductor chips
 chip crisis, 30, 34, 160, 172,
 179–182, 185, 187
 chipification, 34, 172–175, 187

chip production, 35,
 175–179, 182

EU Chips Act, 160,
 183–186, 186n4

Sensing, 1, 2, 10, 22, 23, 33, 49,
 56–58, 60, 63, 72, 76, 182,
 198, 216, 225, 229

Sovereignty, 159–187

Sprenger, Florian, 8, 9, 143,
 144, 152

Srnicek, Nick, 68, 69, 227

Stiftung Neue Verantwortung (SNV),
 177, 180, 181

Strategic Computing Initiative (SCI),
 25, 47, 48

T

Taiwan Semiconductor
 Manufacturing Company
 (TSMC), 3, 34, 160, 172,
 175–181, 185

Technography, 18–19

Testing, 10, 13, 17, 33, 34, 49, 52,
 53, 57, 62, 89, 125, 127, 139,
 197, 198, 204, 205, 209, 218,
 228, 231, 233, 250, 252, 254

Thrun, Sebastian, 26, 84

Tooze, Adam, 168, 169, 187

Traffic cones, 36, 226, 231, 247,
 249–251, 255
 coneheads, 226, 231, 247,
 249–251, 255

Training data(sets), 32, 33, 60,
 67–79, 82–85, 88–90,
 95–121, 256

Turner, Fred, 194, 198, 204

U

Uber ATG, 2, 27, 29–32, 43, 49–56,
 50n1, 56n3, 62, 63, 78, 79,
 81, 105, 127–128, 133–135,
 140, 143, 146, 150, 154, 202,
 233, 242

Urmson, Chris, 26, 27

Urtasun, Raquel, 43, 56n3, 58–63,
 81, 82, 105–108

W

Waabi, 2, 31, 32, 43, 55–57, 56n3,
 60–63, 81, 105

Waymo

- Open Dataset, 27, 32, 68, 69,
 71, 78–89
- Open Dataset
 Challenges, 85, 89

Wellington, Carl, 105–108