



Sensor work: enabling the interoperation of autonomous vehicles

Sam Hind ^{1,*}

¹Digital Humanities, Art History and Cultural Practices, University of Manchester, Manchester, UK

*Corresponding author: Sam Hind. Email: sam.hind@manchester.ac.uk

Abstract

This article examines the “sensor work” carried out in the development of autonomous vehicles which, without sensor data, would not and arguably still do not, have the capacity to decide on where, and how, to drive. I begin by discussing three aspects of sensor technologies considered to be the foundation for sensor work being carried out in autonomous vehicle settings, namely the distribution, processing, and sourcing of sensor technologies and sensor data. The article considers how that much of this sensor work aids not only the operation of autonomous vehicles but also their necessary “interoperation.” In studying four specific sensing methods from an operational perspective, I consider how the interoperation between sensing devices and subsequent algorithmic, object-recognition, and motion planning procedures is fundamental to the development of autonomous vehicles.

Lay Summary

This article examines the work carried out on sensor data in the development of autonomous vehicles. The capture, cleaning, and verification of sensor data collected by different imaging technologies such as cameras, lidar, and radar, is considered essential if vehicle systems are to drive themselves. However, there are many problems that arise in the collection and processing of sensor data that require specialized, and often creative, techniques in order to prepare this data for further use by algorithmic systems within the vehicle itself. The article begins by considering three aspects of sensing technologies that define their use in autonomous driving settings: the distribution of sensors within the vehicle, the processing of sensor data between different components, and the sourcing of hardware that makes it possible for sensing technologies to work in the first place. Through an analysis of how these different aspects and components operate together, the article aims to show how sensor-mediated communication works in an autonomous vehicle setting. In order to show the breadth of this work, I examine four specific techniques being undertaken by machine vision researchers.

Keywords: sensing, autonomous driving, interoperability, sensor work, machine vision, AI

Introduction

Autonomous vehicles are dependent on sensor data, without which they would not—and arguably, still do not—have the capacity to decide on where, and how, to drive. However, to make sensor data useful, a significant amount of work must be performed on it. In this, sensor data must be calibrated, cleaned, plotted, and validated, with unruly, unnecessary, and unreliable data points excluded and erased from further use. This article will consider the various kinds of “sensor work” performed by practitioners in autonomous vehicle settings. In doing so, it will argue that much of this sensor work aids not simply the operation of autonomous vehicles, but their inherent “interoperation.” In this, sensor work is integral to facilitate sensor-mediated communication between sensing systems and subsequent algorithmic decision-making systems within an autonomous vehicle.

I discuss what is meant by interoperability with the help of Adrian Mackenzie and Anna Munster’s (2019) work on “platform seeing.” Mackenzie and Munster discuss how the “massive flows and iterations of images across and within devices, platforms and deep learning models are *plat-format*ted in operation” (Mackenzie & Munster, 2019, p. 9, authors’ emphasis), in which “[s]eeing 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). More

broadly, that adapting to the world of machine vision “involves a redistribution of seeing among classes of experts, technical systems, hardware cloud computing and wireless infrastructures, and among new regulatory frameworks and norms” (McCosker & Wilken, 2020, p. 7). Machine vision demands a certain level, and form, of connectivity: enrolling a vast array of systems to make sense of sensor data.

I will begin this article by discussing three aspects of sensor technologies useful for examining the sensor work being carried out within autonomous vehicle settings. These three aspects concern the distribution of sensors and sensing, the processing of sensor data, and the sourcing of computational hardware required for sensors to work. In introducing these three aspects, I trace a trajectory from smartphones to autonomous vehicles, understanding the latter as extensions of mobile sensing devices, “plat-formatted in operation” (Mackenzie & Munster, 2019, p. 9). It is these aspects that not only set the conditions for sensor-mediated communication, but set sensor-mediated communication apart from other forms of (human) communication as well as wider, and prior, forms of computer-mediated communication.

In the following two sections, I define the concept of “interoperability,” drawing on work in media and communication studies on “operative images” (Farocki, 2004) and “operational data” (Walker Rettburg, 2020), before introducing “operational analysis” (Friedrich & Hoel, 2023), an

approach to studying sensing and algorithmic systems in which tasks, practices, and stages of work along with the so-called operational “pipeline” are foregrounded. Through participant observation and discourse analysis of computer vision conferences and technical literature, I empirically identify four different techniques devised by machine vision researchers to attend to the various problems that emerge in ensuring the interoperability of sensor data with, and within, autonomous vehicle systems. I end the article by discussing the implications of studying such sensor work from an interoperational perspective.

Autonomous vehicles as sensing devices

The starting point for this article is the examination of developments in machine vision that are contributing to sensor-mediated communication, with respect to autonomous vehicles. Sensor-mediated communication is being enabled through parallel developments in multiple fields, but many of these developments are being driven by the financial might of big tech companies and, in particular, their continued investment in mobile devices and smartphones.

Thus, a discussion of sensor-mediated communication with respect to autonomous vehicles must, in some sense, begin with a discussion of the rise of mobile devices as *sensing devices*. Mobile devices are increasingly packed with different kinds of sensors, generally capable of “translat[ing]...stimuli such as light, temperature, speed, and vibration...into electrical resistors and voltage signals” (Gabrys, 2016, p. 8) before being converted into digitally readable form. This is a particularly interesting start point as autonomous vehicles can be understood as scaled-up versions of mobile, sensing devices, rather than just extensions of ordinary vehicles. Whilst this is, in many senses, a product of using sensing devices as a heuristic, it succeeds in drawing out some of the key developmental challenges and tensions that have, up until now, stood in the way of delivering autonomous driving in any form itself.

However, automobiles have long been understood both as enabling communication, and as communicative devices themselves. Wilken and Thomas (2019) have considered the car as a “communication platform,” following the work of Featherstone (2004). Here, Featherstone has argued that the automobile enables multiple kinds of communication: first, from the driver “out through the windscreen, windows and mirrors to the inter-automobile moving figuration of cars” (Featherstone, 2004, p. 8) on the road, second via mediated technological forms, connecting the driver “to distant significant others to help the daily business get done” (Featherstone, 2004, p. 8), and third, whilst “others come in via radio or television, or are physically imported as recordings” (Featherstone, 2004, p. 8). In this, Featherstone contends that the modern automobile “becomes not just a vehicle for independent travel, but a platform for multi-tasking” (Featherstone, 2004, p. 8).

Increasingly, others have considered the car as a form of “mobile spatial media” (Alvarez León, 2019) in itself, enabled by the datafication (Hind, 2021; Martens & Zhao, 2021; Meyers & Van Hoyweghen, 2020) and platformization of the driving experience (Hind & Gekker, 2022; Hind et al., 2022; Steinberg, 2021). Likewise, that “our cars, phones, laptops, Global Positioning System devices, and so on allow for the comprehensive capture of the data trails users leave as they go about the course of their daily lives” (Andrejevic & Burdon,

2015, p. 20) such that cars can be said to play a significant role in the rise of the “sensor society” (Andrejevic & Burdon, 2015, p. 31) altogether, with “the proliferation of sensors” generally leading in “the direction of autonomy” (Andrejevic & Burdon, 2015, p. 31) as autonomous vehicles demonstrate. Here, following the argument above, the car simply becomes an extension of the mobile sensing device, operated through digital software, interconnected through apps and application programming interfaces, generating specific kinds of automotive user data, ordinarily funneled back to car manufacturers, and associated technological partners, such as Google/Alphabet or Intel-owned Mobileye (Pink et al., 2018).

Beyond the car itself, Klein and Selz (2000) have discussed the rise of “cybermediation” in the automotive industry, and the possible changes that “emerging electronic intermediaries” would have on how automobiles were sold by manufacturers and certified dealers. Steinberg’s (2021) recent intervention on Toyota as a precursor to contemporary forms of platform capitalism, considers the ways in which organizational models such as the “stack” or the “intermediary,” have combined with technological innovations such as *kanban* cards and *kaizen* production processes to offer new ways of assembling, marketing, and selling automobiles.

In the following three sub-sections, I trace this trajectory of mobile devices-as-sensing-devices to autonomous vehicles-as-sensing-devices, with respect to the distribution of sensing, the processing of sensor data, and the sourcing of hardware required for sensing.

Sensor distribution

Mackenzie and Munster (2019) discuss how contemporary smartphone cameras operate. Rather than considering them as makers of representational images, they understand them as an “entire sensing ‘platform’ capable of carrying out the distribution and integration of different forms of processing” (Mackenzie & Munster, 2019, p. 14). Following this argument, smartphones do not simply possess a “camera” or a discrete camera module capable of capturing photographs all by itself, but a series of integrated, and connected features each contributing to a different part of the processing of photographs nominally captured by a smartphone’s camera.

In recent years, smartphone cameras have increasingly integrated a range of automated features, alongside increasing the pixel count of the cameras themselves, and the number of cameras and camera types inside the device itself. Whilst the first two generations of the Apple iPhone merely contained a fixed-focus 2.0 megapixel camera, the iPhone 13 contains three cameras (front, wide, ultra-wide) dependent on the same camera system, emphasizing developments in the sensor technology integrated into the device, such as “sensor-shift optical image stabilisation” (Apple, 2022) to offer a better photographic experience. In the case of smartphone cameras, these sensors do “not merely *receive* light but *process* light quantities alongside or in tandem with other information” (Mackenzie & Munster, 2019, p. 14, emphasis added). The importance of this platformed sensing was highlighted recently, when it was reported that “opening the camera in certain apps causes the OIS [optical image stabilisation] motor” in the new Apple iPhone 14 Pro Max to “go haywire” (Hern, 2022), with Apple previously warning users that high-amplitude vibrations (such as those caused by a motorcycle engine) could damage or degrade the phone’s OIS motor.

The story suggested that despite ostensibly being a software feature, the presence of a physical motor, and the complications caused by secondary imaging apps, highlight the distributed nature of sensing in contemporary smartphones. Autonomous vehicles equally rely on the distributed, integrated processing of captured data (Hind, 2022b). Just like smartphones have dedicated cameras for different shooting requirements and sensors that help in specific situations, autonomous vehicles are reliant on multi-sensing systems. In general, three technologies have been used in the development and testing of autonomous vehicles: “cameras” capturing video, lidar devices, and radar.¹ Whilst they can be arranged in many different formations, most if not all manufacturers currently testing autonomous vehicles use a combination of all three, with lidar responsible for the bulk of the sensing work itself. It is through the specific combination of these sensor systems that autonomous vehicles render the outside world knowable, offering a form of ‘machinic sensibility’ (Hong, 2016) that, in the case of lidar, involves material contact between light pulses and physical objects such as trees and road signs. It is through these innumerable points of contact that autonomous vehicles nominally communicate with their surroundings.

On Uber’s infamous autonomous vehicle programme, Uber ATG, sold to Aurora in 2020 (Korosec, 2020), a fleet of modified Volvo XC90s were equipped with cameras, lidar, and radar. GM subsidiary Cruise, responsible for running “robotaxi” services in San Francisco (Marshall, 2022), also utilize all three, including so-called “articulating radars” for greater coverage (Fischer, 2020). Even the electric vehicle manufacturer Tesla, whose CEO Elon Musk has long derided lidar as “a crutch” (Hawkins, 2018), entered into a partnership with lidar manufacturer, Luminar (Hawkins, 2021). This recognition of the need to have a variety of sensors, and sensor types, stretches back to the DARPA Grand Challenges of the mid-2000s, of which the winners, like 2005’s Stanford team, equipped their vehicle Stanley, with five laser sensors (i.e. lidar), a camera, and a radar unit (Thrun et al., 2007). Autonomous vehicles thus are necessarily dependent upon different kinds of sensors, multiple sensing devices, different sensing modes, and different assemblages of sensors in order to distribute the processing of sensor data properly.

Sensor processing

As well as distributing and integrating camera-related functions through the smartphone itself, they are also reliant upon increasing the performance of image signal processors (ISPs) in order to *process* data derived from the camera as well as from other modules within the smartphone. These processes are what Mackenzie and Munster understand as “imaging operations” (Mackenzie & Munster, 2019, p. 14). As they continue, ISPs do not simply process image data, nor do they only process data from one system, but “also receives data from, for example, the gyroscope, which provides image stabilization and combines both signals into one digital image” (Mackenzie & Munster, 2019, p. 15). Apple’s sensor-shift optical image stabilization, for instance, uses information generated by the iPhone’s in-built gyroscope to move, or shift, the image sensor itself, via actuators, making up to 5,000 adjustments per second (Hristov, 2021). Hence, why high-amplitude vibrations may disrupt such a process, dependent on the making of high-speed physical corrections.

The transformation of smartphones into sensing devices, however, has only been made possible through the development of graphics processing units (GPUs). As Mackenzie and Munster suggest, “GPU architecture, the silicon substrate of millions of first-person standpoint 3D action games, with their pursuit of detailed and fluidly mobile game physics, has developed to render images aggregately computable through massive calculative parallelism” (Mackenzie & Munster, 2019, p. 17). This “calculative parallelism” enables “vast numbers of discrete arithmetic operations” to be “carried out in parallel lanes” (Mackenzie & Munster, 2019, p. 17) to generate images. It is this calculative parallelism that can be witnessed in the sensor work detailed herein, as autonomous vehicle practitioners likewise grapple with how to handle huge volumes of sensor data. Indeed, as Mackenzie and Munster reiterate, smartphone ISPs are merely a “downsized iteration of...image recognition processors for autonomous vehicles” (Mackenzie & Munster, 2019, p. 16), such that for big tech firms like Apple and Google, they function as the computational foundations of scalar operations, like autonomous vehicle projects.

As Wilken and Thomas (2019) discuss, autonomous vehicles generate huge volumes of data, requiring innovative techniques to formalize, standardize, and prioritize sensor-mediated communication. For instance, in how lidar systems can be “calibrated to prioritize either the strongest or last distance point recorded” (Hind, 2022b, p. 69) in order to discount certain environmental conditions that may inhibit sensor-mediated communication, like fog or dust. Likewise, that under certain test conditions some sensing systems might be deactivated completely in order to avoid communication conflicts where interoperability between such systems has not yet been achieved, and certain vehicle components (such as a braking system) have not been programmed to prioritize commands sent from elsewhere. In cases such as the Uber ATG crash in Tempe, Arizona in 2018, the Volvo XC90 vehicle involved was equipped with an integrated advanced driver assistance system (ADAS), disabled by Uber technicians to avoid communication conflicts with in-house sensing systems being developed (Hind, 2022b).

Sensor sourcing

Another important factor is the *sourcing* and assembly of semiconductor chips, now otherwise known as central processing units, key components that combine with GPUs to offer visual processing (Forelle, 2022). As the semiconductor “chip crisis” of 2020–2022 developed, as a result of the global Covid-19 pandemic, it became apparent that the sourcing of semiconductor chips was becoming increasingly difficult. Although the chips required for automobiles are not the same kinds of chips for smartphones and other mobile devices, chip fabrication facilities or “chipfabs” cannot easily be converted to produce different chip sizes. As a result, there became a significant backlog of semiconductor chip orders, such that car manufacturers were faced with considerable production delays and vehicle models essentially being “sold out” for the year. The global automotive industry, thus, made 7.7 million fewer cars in 2021 than 2020 (Ting-Fang & Li, 2022). As well as political decisions made by the European Union (EU) and USA to address the mid- to long-term supply-chain issues, embodied in the tabled EU Chips Act and the US Chips and Science Act, manufacturers resorted to downgrading vehicle models previously reliant upon unavailable chips (Knight,

2022; Szymkowski, 2021). These aspects are important to consider as they modulate the political (legislation), economic (manufacture), and technical (design) conditions for the adoption of sensor-mediated communication, and the acceleration of the ‘sensor society’ (Andrejevic & Burdon, 2015), as mentioned previously.

Compounding these problems was the fact that semiconductor chip firms such as the Taiwan Semiconductor Manufacturing Company (TSMC), have themselves been suffering from supply issues (Ting-Fang & Li, 2022). Companies such as Screen, a Japanese firm specializing in chemical cleaning equipment essential to the chip fabrication process, were having to inform their clients that “valves, tubes, pumps and containers made of special plastics” (Ting-Fang & Li, 2022) for such equipment were difficult to source. As a result of these “cascading” problems (Ting-Fang & Li, 2022), chip firms had no choice but to extend already delayed delivery dates for clients, such as automotive manufacturers. In light of cascading supply-chain sourcing issues of everything from raw materials to pipes, tubes, and pumps, political attempts to “onshore” or regionalize semiconductor supply-chains have faced significant challenges. Thus, Intel’s decision to build a new chipfab in Magdeburg, Germany (Tagesschau, 2022), in light of the EU Chips Act, and the EU’s desire to reduce its dependence on “third-country suppliers” such as Taiwan and South Korea (European Commission, 2022), will not necessarily resolve these “upstream” supply-chain issues, more rudimentary, but no less troublesome, in nature.

Sensor work/defining interoperability

In order to understand the work that is carried out in relation to the distribution, processing, and sourcing of sensor technologies and sensor data in the development of autonomous vehicles, I turn to the concept of “interoperability.” I define interoperability as the transmission of sensor data from one (sensing) system to other, connected systems deemed necessary for subsequent decision-making processes. I understand interoperability as more than “interconnectivity” in that (sensor) data must be “formatted” (Volmar et al., 2020) for more than one of these systems, in order to flow through them. Here, interoperability is also different from “co-operation” in that respective sensing and decision-making systems do not necessarily work together, or co-operate, to achieve a mutual aim, but interoperate to achieve specific modular, or parallel, goals such as detecting 3D objects or processing video frames. With this, these tasks are merely constituent parts of a continual, operational chain.

As Volmar et al. suggest, whilst the term “format” describes “structural or programmatic relationships between individual elements and their organizational logic” (Volmar et al., 2020, p. 8), something like interoperability or interoperationality concerns the quality or form of the contents subject to such a structural or programmatic relationship. Thus, the formatting of individual elements such as sensor data is important, but for this article, only insofar as it offers an understanding of how such a formatting aids or enables interoperability.

This article’s interest in interoperability is an extension of existing work on “operative images” (Distelmayer, 2018; Farocki, 2004; Hoel, 2018) and “operational data” (Walker Rettburg, 2020). In this work, the notion of “operative” images or “operational” data are derived through a

distinction with representational images and/or data, which are not utilized for automated purposes, nor principally designed to be viewed or read by human actors, but by machinic systems (Farocki use the example of guided missiles). As Walker Rettburg (2020, p. 9) considers, operational data are data that is “algorithmically processed...with little human involvement” with “no need for human-readable representations.” Whilst, in a heuristic sense, it is necessary to draw a distinction between human-readability on the one hand (representation) and machine-readability on the other (operational), much human work is still required to ensure the ongoing operability of machine-readable data, not least in respect to the sensor data captured by autonomous vehicles and their sensing systems.

Another critical distinction with work on operative images/operational data is that, as I want to argue here, interoperationality precedes operability, rather than vice versa. If one understands an operative image or operational data as having an operative/operational being or existence, this is definable in relation to their operability within a specific, holistic system. However, within the context of autonomous driving, there is no such single, specific, holistic system; only distinct, integrated systems for different tasks (sensing, processing, control, etc.), such that it is inaccurate to state that the sensor data flowing through, and between, such systems merely “operate.” Instead, the interoperability and interoperationality of each unit of sensor data must be ensured before any one system “operates.” Here, rather than interoperability following operability, operability is dependent firstly on interoperability.

This continual, operational chain or “pipeline” along which sensor data must flow, is critical to the eventual “success” of autonomous driving—at least according to an analysis of recent work in the industry. Moreover, that the history of the development of autonomous vehicles, stretching back to Stanley, Uber ATG, and Cruise, suggests that autonomous driving demands *thinking* and *acting* in an interoperational fashion. As the section before introduced, sensor work in the development of autonomous vehicles is dependent on three aspects: distributing sensor devices, processing large volumes of sensor data from these varied sources, and securitizing the sourcing of semiconductors and related components underpinning both. Without each, the *technical* interoperationality of autonomous vehicles hits a decidable snag.

Whilst this observation might appear obvious—that sensing systems must be integrated with other systems—work to develop them routinely takes a modular, and often sequential, form, as leading autonomous vehicle engineer Urtasun (2021) has argued. In other words, that separate engineering teams in any autonomous vehicle firm, first and foremost, are responsible for their own systems. Any integration with other systems usually comes second. As a result, any technical solutions or methods devised to resolve problems related to the operation of these discrete systems, typically involves ‘developing more and more modules’ (Urtasun, 2021). Any resultant issues concerned with integration and interoperation are therefore considered as secondary, and usually therefore lesser, problems. This is despite the critical importance of having such systems work inter-operationally. Conceiving interoperability differently, therefore, requires a fundamental re-organization of both workflow and team composition, beyond a “traditional engineering stack” (Zeng et al., 2021, p. 1) that favors discrete, subdivided, tasks.

Thus, at this current stage in the development of autonomous vehicles, it is important to understand what I call here the “sensor work” being undertaken to provisionally deliver autonomous driving. Put otherwise, what is sensor work and how is it organized to enable interoperability? This work, as I have suggested already, involves experimentation with, and in, machine vision in order to smooth the interoperability of sensor-mediated communication between sensing systems. In short, that there are parallel innovations in relation to sensor distribution, processing, and computation. By focusing on this practical sensor work, we can begin to understand how sensor-mediated communication is enabled in the realm of autonomous driving, and begin to gain a greater sense of how the work around it is structured.

This is an important theoretical step, as the powers of sensing technologies do not singularly reside in any discrete, nominal sensing technology, device, or even system, but are imbued at the various stages or moments along the operational chain/pipeline, at which different components are *theoretically* although not always *actually* “plugged into” each other. It is at these moments that sensor work is done, and through which interoperability is nominally achieved, and without which the entire sensing assemblage either functions sub-optimally or breaks down entirely.

Methods

Methodologically, the article offers an “operational analysis” of sensing and sensor work, drawing on Bucher (2018), Marres (2020), Friedrich and Hoel (2023), and Rieder and Skop (2021), who variously approach the study of algorithmic systems from an operational, and situational, perspective. Here, “operation” refers to the technical operation of a computational system in question, whether the integrated algorithmic processes underpinning social media platforms (as in Bucher), robotic stereotactic radiosurgery systems (as in Friedrich and Hoel), automated content moderation tools (as in Rieder and Skop), or autonomous vehicle systems.

An operational approach avoids a macro- or media-centric perspective on such phenomena, which might seek to establish, or at least proceed to investigate, the enduring properties of the systems in question. In such a case, social media platforms, robotic radiosurgery systems, automated content moderation tools, or autonomous vehicle systems might come to be understood as such through a search for, and definition of, specific features (an algorithmically sorted newsfeed, a banned word detector) or components (a motion synchronization device, a lidar unit) that comprise them. Falling short, perhaps, of a fully typological approach, such a perspective nonetheless results in the categorization of various kinds of technical systems into what they are provisionally *meant to do* or *said* to be able to do at the expense of their understanding as situated technologies.

At the other end of the spectrum, an operational approach also avoids a strictly micro- or localized analysis that shears specific performance of such systems either from generalizable or categorizable properties, or from the iterative, relatable work practices that enroll and span across them, instead seeking to establish the originality of each instance of use. In other words, a localized (or perhaps “hyper-localized”) account might avoid making comparisons between how robotic radiosurgery is performed in one hospital setting as opposed to

another, or how one content moderation tool is deployed in a newsroom compared to a social media company.

Instead, operational analysis adopts a meso- or “middle-range” (Friedrich & Hoel, 2023, p. 52) approach designed to follow the operational “task...to be performed” (Friedrich & Hoel, 2023, p. 52). In this article, this means attending to the “fabrics” (Rieder & Skop, 2021, p. 2) of sensing and sensor work, as a “variety of actors, sites, processes, [and] technologies” (Rieder & Skop, 2021, p. 13) are weaved together to achieve certain operational goals. For Marres, this entails the recognition of “the dynamic nature of situations” (Marres, 2020, p. 7) and, consequently, the “unfolding of situations in computational settings” (Marres, 2020, p. 7).

More systematically, this means identifying a “task of interest” (Friedrich & Hoel, 2023, p. 59) such as the treatment of a tumor, or the moderation of comments on a news article, before delineating the so-called “operative moments” (Friedrich & Hoel, 2023, p. 60) occurring in relation to the task, such as the correct positioning of a patient, or the flagging of banned words. An operational analysis, therefore, considers the socio-technical operation of particular systems through the lens of “tasks” and the necessary multistage processes or “operative moments” through which systems are put to work.

I draw on four processes encountered in research into machine vision in an autonomous vehicle context that are best explicated through such a “middle-range” approach. These include: lidar point attenuation, 3D object detection, streaming processing optimization, and depth sensor processing.

This research included participant observation of a virtual summit on computer vision (“Machines Can See”) hosted in June 2020 and a virtual workshop on autonomous driving hosted in June 2021 (“Workshop on Autonomous Driving” or WAD).² International summits, conferences, and workshops concerning the technical aspects of sensor work are key occasions in which cutting-edge methods, techniques, and approaches are shared with those working in related fields of computer vision and deep learning. Many of these events, such as the Computer Vision and Pattern Recognition (CVPR) conference (which hosted WAD), are principally organized by academics in the field, on behalf of US organizations such as the Institute of Electrical and Electronics Engineers and Computer Vision Foundation (CVF).

However, they also typically have connections to major commercial firms with an interest in computer vision and machine learning. Whilst CVF, for example, describes itself as a “non-profit organization that fosters and supports research in all aspects of computer vision” (CVF, 2022), it is also funded by four “diamond sponsors”: Amazon, Microsoft, Google, and Facebook. Similarly, CVPR has a substantial number of “platinum sponsors” including Amazon, Apple, Qualcomm, Tencent, and Datagen, as well as autonomous vehicle firms Argo, Cruise, and Waymo. Argo and Waymo are also recurring (2020–2023) participants in WAD, offering keynote speakers, and prize monies to entrants (WAD, 2022, 2023). As a result, work presented at such conferences on image recognition, data segmentation, and visual processing is of considerable commercial interest.

Yet, on their own these events do not necessarily capture the liveness, and diversity, of approaches to computer vision and machine learning. Accordingly, it is important to be able to trace and connect the circulation of key technical papers published by those within related fields. This orientation towards the “social life” of machine learning methods (Savage, 2013)

and associated sensor work was attained through the organization of two “hands-on” workshops. The first, “Making Sense of Sensor Data,” provided the opportunity to explore how sensor data is ordinarily used to develop training datasets for machine learning purposes, such as with the KITTI dataset (Geiger et al., 2012) or Waymo’s Open Dataset (Waymo, 2019). The second workshop, “Taking up the Challenge,” offered the opportunity to examine how these training datasets are used in so-called machine vision “challenges,” such as in Waymo’s Open Dataset Challenges (Waymo, 2020).³

Convention within the world of computer vision and machine learning means that novel methods or techniques are uploaded onto open-access paper repositories such as arXiv and popular code repositories such as Github. Through these common practices, users of both platforms can not only read about specific methods or techniques, but also potentially use them as scaffolding or “backbones” for additional work—something Rieder and Skop refer to as a “cooperative, multipolar model” (Rieder & Skop, 2021, p. 10). Whilst this article does not offer a systematic analysis of the web of connections each technique ordinarily offers, these features have been used to gauge the impact each published method has had, and the prior work it is dependent on, such as the datasets used to train particular machine learning models. The arXiv submission for the Mask R-CNN method used for object instance segmentation, cited by 20,844 other academic articles, used both the COCO and Cityscapes datasets, two significant image datasets used for training object recognition/segmentation models (He et al., 2018).

Both aspects (industry conferences, self-organized workshops) that I have introduced here build on extant knowledge of where, and how, machine vision research is taking place, both commercially and academically, arising from past work on autonomous vehicles (Hind, 2019, 2022a, 2022b).

Sensor work in action

An examination of these sensor processes not only intends to surface the sensor work needed to make autonomous vehicles themselves “work,” but also the range of possible *strategies* devised, and devisable, by practitioners to deal with the necessary interoperability of autonomous vehicles. As the final section of the article will examine, these strategies involve a great degree of invention and inventiveness to develop possible, plausible, and workable, methods for various sensor-related issues, such as removing the interference of “spurious” objects like rain droplets, or tackling “stale” video frames. In this, sensor-mediated communication within autonomous vehicles is only made possible by such work, without which autonomous vehicles would cease to offer “autonomous” decision making and movement. I proceed by considering each of the processes in turn as they appear chronologically in the operational pipeline, from initial sensing to the design of the so-called “bounding boxes.”

Lidar point attenuation

The first sensor process involves the attenuation, or amelioration, of sensor data “noise” (Espineira et al., 2021). Lidar point attenuation is an important process because of the extent to which lidar points, as rays of light, are distorted by other objects. The most problematic of these other objects, yet perhaps also the least perceptible, are raindrops and dust particles. Whilst both phenomena invariably play a part in how

any autonomous vehicle proceeds through an environment, the material presence of rain droplets and dust particles close to a lidar unit has the potential to significantly disrupt sensor processes. As Espineira et al. (2021, p. 8) write, “if a lidar beam intersects with a raindrop at a short distance from the transmitter, the raindrop can reflect enough of the beam back to the receiver such that it is detected as an object.”

Likewise, that raindrop interference can also degrade the strength and the so-called “range performance” (Espineira et al., 2021, p. 8) of a lidar device itself. Thus, if lidar data are to be made useful, engineers must be attuned to the variable properties of these troublesome atmospheric agents. Without knowledge of exactly how raindrops and dust particles affect the operation of lidar devices, researchers are likely to assume minimal mitigation strategies are necessary. Whilst dust particles may not present quite as frequent a problem for autonomous vehicles typically driving on tarmac roads, tackling them first emerged as a problem during the original DARPA Grand Challenges, set on the dusty, desert roads of Nevada, USA (Buehler et al., 2007).

Espineira et al.’s solution involves the modelling of rain droplet particles within a virtual environment. What is important here is that such a “probabilistic rain model” (Espineira et al., 2021, p. 8) must be made interoperable with any real-world, real-time lidar device, such that it is capable of ingesting a “pointcloud affected by the rain in the same format and data-rate of [a] pointcloud generated by a real lidar sensor” (Espineira et al., 2021, p. 7). In order to do so, the researchers used the Unreal Gaming Engine, software typically used to design virtual gaming environments, but increasingly used within autonomous vehicle development to design simulations (Steinhoff, 2022). The Unreal Gaming Engine allowed them to develop a real-time rain model through which they could modify raindrop size (between 0.5 mm and 6 mm) and rainfall intensity (e.g. between 10 mm/hr and 50 mm/hr, for moderate and heavy rainfall). By running a simulated, realistic lidar point cloud within the game engine, in concert with the rain model, the researchers were, in principle, able to generate different scenarios in which rain droplets variously interfered with lidar point distribution.

As Espineira et al. (2021, p. 8) state, this “opens the possibility to study the combined real degradation of visual and lidar sensors in the case of rainfall of different intensities [as well as] the possibility to vary the detector threshold...tailored to the performance of a specific commercial lidar.” In effect, that their virtual environment could be used to modulate, or fine-tune, the sensitivity of any particular lidar sensor, to enable it to ignore, or more specifically to “discount,” rain droplets at various distances. Moreover, that such fine-tuning could easily be performed with mixed engineering teams, each of those with an express interest in ensuring the lidar data entering subsequent systems is clean and useful.

Dealing with the noise generated through such sensing processes is of critical importance, as any sensor data that subsequently passes along the pipeline to be acted upon by object-recognition algorithms might, without such attenuation work, result in the autonomous vehicle responding to so-called “hallucinations” (Kayhan et al., 2021), objects that are not really there, or do not pose as big a problem as their recorded “object-ness” supposes.⁴ Thus, we can understand this sensor work as ideally exhibiting a level of interoperability between simulated rain model and real-world scenario (if done

correctly), that ultimately should *refine* interoperability between the perception and recognition stages of the decision-making process, as techniques for discounting rain droplets, dust particles, and similar objects are developed.

3D object detection

The second sensor process concerns the ability to detect and orient vehicles within 3D space (Hu et al., 2020). In short, to improve the level of visibility enabled through the chosen method of sensing. Rather than decide to use another type of sensor, Hu et al. stick with lidar, choosing to engage with the operational issues it throws up. That main issue, as the authors write, is that lidar points effectively destroy or otherwise elide data on phenomena behind each lidar point captured. This is because a lidar point will be returned when it hits an object. Any secondary object hidden behind this initial object, therefore, is not captured at all, essentially destroyed within the dataset being generated. As they contend, “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” (Hu et al., 2020, p. 2). 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), as some objects physically represented in space cannot, and will not, be captured each time a lidar scan is performed. Naturally, this could represent a problem for any autonomous vehicle, hampered by lidar’s “occlusionary” capacities.

Their solution is a technique referred to as “raycasting.” Ordinarily when a lidar point is recorded it generates a specific coordinate in 3D space, the exact point at which an object was hit, and the lidar point returned. Repeated many thousands of times for every lidar “sweep” and one builds up a series of positively recorded coordinates. However, the so-called “freospace” in between each object and the lidar unit is not typically recorded. Hu et al. propose a method for generating a “3D voxel grid” in which each coordinate is recorded as either “occupied, free, or unknown” (Hu et al., 2020, p. 4), thereby allowing each point within 3D space to be provided with a value, regardless of whether a lidar point has been generated. They refer to this as a “visibility volume” (Hu et al., 2020, p. 4), and can be thought of as a technique for turning the 2.5D of lidar into fully 3D data.

Here, what becomes more evident is how the re-formatting of lidar data into “full” 3D form, with the help of an intermediary (the 3D voxel grid), is designed to yield a greater level of interoperability between lidar input and a desired feature map output, key to the “diagrammatic abstractions” (Mackenzie, 2017, p. 55) made through, and required by, machine learning. Without the 3D voxel grid, so the researchers would argue, the lidar input offers an impoverished view of the world, unable to properly record all manner of occluded objects. What this stated technique offers, therefore, is a way of enhancing or *augmenting* the perceptive qualities of lidar, such that richer, fuller, and deeper snapshots of the immediate environment can be captured.

Streaming processing optimization

The third sensor process entails the speeding-up, and economization, of video frame processing (Li et al., 2020). In principle, “streaming processing optimization,” a term coined by researchers at the Argo AI Center for Autonomous Vehicle Research (Carnegie Mellon University, Pittsburgh, USA), is designed to tackle one particularly troublesome issue in

vehicle perception. This is the trade-off between accurate and quick image understanding, where accuracy is defined by achieving a certain threshold of objects correctly identified and categorized, and speed is defined by completing an image understanding process in advance of subsequent phases in the operational pipeline, such as motion forecasting and planning. An image understanding process that takes too long to complete would result in objects within such images (namely, other road users) not being properly identified/classified.

For example, if a cyclist was improperly categorized as an ordinary vehicle and upon crossing the autonomous vehicle was still deemed to be moving at the (faster) speed of a vehicle, rather than a bicycle. Consequently, the cyclist might be at a higher risk of being hit by the autonomous vehicle, believing the cyclist was traveling faster than they were. Thus, optimizing the processing of such image recognition tasks is of critical importance.

The solution offered by Li et al. (2020, p. 1) is something they call “dynamic scheduling.” Counterintuitively, the algorithmic object-recognition system addresses any potential latency problem between perception system and world by “sitting idle and ‘doing nothing’” (Li et al., 2020, p. 1) rather than tackle so-called “stale frames” (Li et al., 2020, p. 20), snapshots of the immediate real world, now in the past. Thus, rather than have an object-recognition system sequentially line-up video frames to process, fully completing one frame before moving onto the next, if the current state of the real-world has progressed beyond the workflow of the system (say, if the vehicle has increased its speed, thereby increasing the necessary speed at which frames must be processed, and objects identified/classified), it simply waits until it is once again able to process the current situation.

As Li et al. (2020, p. 1) consider, a “crucial quantity governing the responsiveness of [an autonomous] agent is its reaction time.” Providing it the opportunity to wait, rather than processing stale frames, so their argument goes, gifts the system the opportunity to react quicker to current situations as they emerge. Thus, dynamic scheduling smooths the interoperation between video input and classification outputs, meaning that critical processing capacity is not wasted by parsing “useless” frames. Arguably in the absence of such a method engineers would face a greater hurdle in processing images both quickly and accurately. What this technique arguably offers is a recalculation of the necessary *pacing* of interoperability—not that video inputs and classification outputs need to interoperate, but that interoperation proceeds situationally, with the system cognizant of the value of the frame to be processed, rather than ignorant of its irrelevance to the current situation.

Depth sensor processing

The final sensor process is the generation of 3D “bounding boxes” from monocular RGB images (Gähler et al., 2020). This is a process developed by researchers associated with the popular Cityscapes dataset, as mentioned earlier, often used by autonomous vehicle projects for training object-recognition models. The technique they have developed is designed to avoid “sensor setup, calibration and synchronization” (Gähler et al., 2020, p. 2) issues commonly encountered when using lidar as part of a multi-sensing system. One issue they identify is the “sparsity of lidar measurements, especially for distant objects” (Gähler et al., 2020, p. 5). Another is that 3D bounding boxes derived from lidar data

can “result in imprecise reprojections into...RGB images” (Gähler et al., 2020, p. 3), as synchronization issues between individual lidar sensors might result in being unable to correctly capture fast-moving objects, especially those closest to the autonomous vehicle itself. Invariably some bounding boxes, although they might accurately contain the extent of the object in question, also do not take road curvature into account, meaning that pitch and roll annotations are improperly defined. For cities or areas with steep inclines or undulating terrain, this phenomenon has the potential to be both significant and frequent, especially for vehicles nudging out of steep side-roads, as Gähler et al. show.

Their solution is to develop a method to generate 3D bounding boxes using only 2D images. Here, rather than trying to correct the various calibration and synchronization issues with lidar mentioned above, post facto, depth sensor processing generates 3D bounding boxes from 2D images captured by camera alone. Whilst Gähler et al. manage to avoid any such synchronization issues between sensing systems and types of sensor data (lidar, camera), there are still necessary tradeoffs. As cameras do not have the possibility to “see” through physical objects, depth sensor processing must involve an evaluation of whether an object is “occluded” (i.e. obscured) or “truncated.” All further objects proceed to be classified into one of 24 further classes, corresponding to the general dimensions of different vehicle types such as a sedan or large van.

In developing this alternative, so the authors argue, there is a greater *synchronization*, or interoperability, between raw sensor data and the algorithmic annotation process, that does not require integration between different sensing modes. Here, lidar is itself understood as an impediment to interoperability, especially as its use ultimately requires the translation of 3D point clouds into 2D outputs. Yet, as this approach shows, interoperability does not produce binary states, where interoperability is either achieved or it is not. What is lost through the above, is the depth of sensing offered by lidar, unable to be matched by any camera. In strengthening the interoperability between image inputs and bounding box outputs further down the pipeline, the method only weakens the perceptive capacities of the vehicle in the first place.

Conclusion

In this article, I have sought to provide an account of the so-called “sensor work” being carried out in the development of autonomous vehicles. Necessarily incremental, such work deals with the knotty question of “interoperability”: how sensor data must not only be captured, but also cleaned, corrected, formatted, and re-presented in order to pass further along the operational “pipeline,” to be subsequently processed by object-recognition algorithms and motion planning modules in any autonomous vehicle system.

Whilst much of the technical work discussed above does not refer to interoperationality by name, all of this work proceeds on the basis that the sensor data being used has “somewhere else to go.” But that, critically, before it is sent along the pipeline that certain sensor-related issues must be tackled in advance. However, none of the proposed techniques necessarily provide ultimate solutions, nor cast-iron guarantees, that using such methods will entirely eradicate the various problems inherent to the distributive, processual, and computational capacities of sensor technologies and sensor

data. Instead, they all proceed on the basis that each problem they orientate their technique towards is conditionally resolvable, conditional on an array of factors such as the distributive arrangement of sensors, the speed of processors, or the overall computational capacity.

For the question of sensor-mediated communication these are, hopefully, valuable observations. Both the “rise of sensors” and the “explosion of sensor-derived data” (Andrejevic & Burdon, 2015, p. 25) have yielded ever-more complex infrastructural, organizational, and operational arrangements in order to capture, process, order, and value myriad forms of sensor-mediated communications. That much of this takes place out of sight, beyond the immediate purview of ordinary users and operators of such technical systems, is a sign that the landscape of both computer-mediated communication and human-computer interaction is rapidly shifting, such that sensor-mediated communication can now be considered a default mode in many settings.

In many of the techniques discussed here, there is an explicit acknowledgement that ideal conditions do not exist, and that each method must therefore take account of the specificity of each problem, as well as the technical set-up of each autonomous vehicle system. In other words, that the researchers broadly acknowledge, and work on the basis, that their methods have a positionality to them, necessarily situated and “biased” to the conditions of their making, like all projects involving algorithms and machine learning (Jaton, 2021).

The principle theoretical foundation for the article has therefore been interoperability. What I have sought to argue within is that to understand sensor-mediated communication one must necessarily begin from the interoperation between, rather than the operation of, specific sensing systems and connected, interpretive, algorithmic systems. In respect to autonomous driving, as Zeng et al. (2021, p. 1) write: “most...companies have large engineering teams working on each sub-problem in isolation,” each working to a sub-system solution. Yet, approaching sensor work in a modular fashion inevitably generates problems for engineers further down the (pipe)line, where an “advance in one sub-system does not easily translate to an overall system performance improvement” (Zeng et al., 2021, p. 1).

If we are to properly understand how sensor-mediated communication is enabled as it proliferates throughout the world, in new sensing and algorithmic contexts, and through novel sensing situations, then it would seem sensible to pay attention to where, and how, work is being done to enable interoperability—and where the organizational, or technical, hurdles to ensuring it may lie.

Data availability

The data underlying this article will be shared on reasonable request to the corresponding author.

Funding

The research for this article was funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation)—Project-ID 262513311—SFB1187 Media of Cooperation.

Conflicts of interest: None declared.

Notes

1. Whilst there seems to be a sharp distinction between “cameras” and “sensors” in the smartphone world, such a distinction is harder to make in the world of autonomous vehicles. Even though vehicles may be equipped with video recorders to capture footage, they are still ordinarily referred to as “cameras” rather than recorders. Indeed, whilst lidar and radar are ordinarily referred to as “sensors,” use of terms like “see” and “sight” colloquially convey their qualities as modes of sensing, despite not offering the same kind of sight as cameras or video recorders.
2. The “Machines Can See” summit was held online from June 8 to June 10, 2020, and the “Workshop on Autonomous Driving (WAD)” was held online on June 20, 2021.
3. The workshop “Making Sense of Sensor Data” was held online from November 8 to November 10, 2021 and organized by members of the A03 Navigation in Online/Offline Spaces project of SFB1187 Media of Cooperation at the University of Siegen, Germany. The workshop “Taking up the Challenge” was held from July 14 to July 15, 2022 and also organized by members of the A03 Navigation in Online/Offline Spaces project of SFB1187 Media of Cooperation at the University of Siegen. I thank Max Kanderske and Fernando van der Vlist for co-organization of these events.
4. I thank Susanne Förster for first alerting me to the concept of hallucinations in machine vision work.

References

- Alvarez León, L. F. (2019). How cars became mobile spatial media: A geographical political economy of on-board navigation. *Mobile Media & Communication*, 7(3), 362–379. <https://doi.org/10.1177/2050157919826356>
- Andrejevic, M., & Burdon, M. (2015). Defining the sensor society. *Television & New Media*, 16(1), 19–36. <https://doi.org/10.1177/1527476414541552>
- Apple (2022). *iPhone 13*. Apple. <https://www.apple.com/uk/iphone-13/>
- 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. (2007). *The 2005 DARPA grand challenge: The great robot race*. Springer.
- Computer Vision Foundation (2022). *The computer vision foundation*. CVF. <https://www.thecvf.com/>
- Distelmayer, J. (2018). Carrying computerization: Interfaces, operations, representations. 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>
- European Commission (2022). *European Chips Act: Questions and answers*. European Commission. https://ec.europa.eu/commission/presscorner/detail/en/QANDA_22_730
- Espineira, J.P., Robinson, J., Groenewald, J., Hung Chan, P., & Donzella, V. (2021). Realistic LiDAR with noise model for real-time testing of automated vehicles in a virtual environment. *IEEE Sensors Journal*, 21(8), 9919–9926. <https://doi.org/10.1109/JSEN.2021.3059310>
- Farocki, H. (2004). Phantom images. *Public*, 29, 12–22.
- Featherstone, M. (2004). Automobilities: An introduction. *Theory, Culture & Society*, 21(4–5), 1–24. <https://doi.org/10.1177/0263276404046058>
- Fischer, J. M. (2020, July 1). The decision behind using articulating sensors on Cruise AVs. *Cruise*. <https://medium.com/cruise/cruise-embedded-systems-articulating-radars-7cae24642930>
- 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>
- 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>
- Gabrys, J. (2016). *Program earth: Environmental sensing technology and the making of a computational planet*. University of Minnesota Press.
- Gährlert, N., Jourdan, N., Cordts, M., Franke, U., & Denzler, J. (2020). *Cityscapes 3D: Dataset and benchmark for 9 DoF vehicle detection*. ArXiv. <https://doi.org/10.48550/arXiv.2006.07864>
- Geiger, A., Lenz, P., & Urtasun, R. (2012). Are we ready for autonomous driving? The KITTI Vision Benchmark Suite. In *2012 IEEE conference on computer vision and pattern recognition*, Providence, RI, June 16–21, 2012, pp. 3354–3361. <https://doi.org/10.1109/CVPR.2012.6248074>
- Hawkins, A. J. (2018). *Elon Musk still doesn't think LIDAR is necessary for fully driverless cars*. The Verge. <https://www.theverge.com/2018/2/7/16988628/elon-musk-lidar-self-driving-car-tesla>
- Hawkins, A. J. (2021). *Elon Musk called lidar a “crutch,” but now Tesla is reportedly testing Luminar's laser sensors*. The Verge. <https://www.theverge.com/2021/5/24/22451404/tesla-luminar-lidar-elon-musk-autonomous-vehicles>
- He, K., Gkioxari, G., Dollár, P., & Girshick, R. (2018). *Mask R-CNN*. ArXiv. <https://doi.org/10.48550/arXiv.1703.06870>
- Hern, A. (2022, September 19). Bug in iPhone 14 Pro Max causes camera to physically fail, users say. *The Guardian*. <https://www.theguardian.com/technology/2022/sep/19/iphone-14-pro-max-camera-optical-image-stabilisation-motor>
- 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. (2022a). Making decisions: The normal interventions of Nissan ‘mobility managers’. *Mobilities*, 17(4), 467–483. <https://doi.org/10.1080/17450101.2021.1988682>
- Hind, S. (2022b). 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://www.aup.nl/en/book/9789463727037/seeing-the-city-digitally>
- 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>
- 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>
- Hong, S-H. (2016). Data’s intimacy: Machinic sensibility and the quantified self. *Communication*, 5(3), 1–36. <https://doi.org/10.7275/R5CF9N15>
- Hristov, V. (2021, September 2). What is sensor shift stabilization and is it better than optical image stabilization (OIS)? *Phone Arena*. https://www.phonearena.com/news/What-is-sensor-shift-stabilization-vs-OIS_id133572
- Hu, P., Ziglar, J., Held, D., & Ramanan, D. (2020). *What you see is what you get: Exploiting visibility for 3D object detection*. ArXiv. <https://doi.org/10.48550/arXiv.1912.04986>
- Jaton, F. (2021). *The constitution of algorithms: Ground-truthing, programming, formulating*. MIT Press. <https://doi.org/10.7551/mitpress/12517.001.0001>
- Kayhan, O. S., Vredbregt, B., & van Gemert, J. C. (2021). Hallucination in object detection: A study in visual part verification. ArXiv. <https://doi.org/10.48550/arXiv.2106.02523>

- Klein, S., & Selz, D. (2000). Cybermediation in auto distribution: Channel dynamics and conflicts. *Journal of Computer-Mediated Communication*, 5(3). <https://doi.org/10.1111/j.1083-6101.2000.tb00347.x>
- Knight, S. (2022, May 30). Companies are hacking their way around the chip shortage. *Wired*. <https://www.wired.com/story/chip-shortage-hacks/>
- 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/>
- Li, M., Wang, Y.-X., & Ramanan, D. (2020). *Towards streaming perception*. ArXiv. <https://doi.org/10.48550/arXiv.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 invisibilities. *Theory, Culture & Society*, 36(5), 3–22. <https://doi.org/10.1177/0263276419847508>
- McCosker, A., & Wilken, R. (2020). *Automating vision: The social impact of the new camera consciousness*. Routledge.
- 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>
- 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/>
- Martens, B., & Zhao, B. (2021). Data access and regime competition: A case study of car data sharing in China. *Big Data & Society*, 8(2), 1–11. <https://doi.org/10.1177/20539517211046374>
- Meyers, G., & Van Hoyweghen, I. (2020). "Happy failures": Experimentation with behaviour-based personalisation in car insurance. *Big Data & Society*, 7(1), 1–14. <https://doi.org/10.1177/2053951720914650>
- Pink, S., Fors, V., & Glöss, M. (2018) The contingent futures of the mobile present: Automation as possibility. *Mobilities*, 13(5), 615–632. <https://doi.org/10.1080/17450101.2018.1436672>
- 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>
- Savage, M. (2013). The "social life of methods": A critical introduction. *Theory, Culture & Society*, 30(4), 3–21. <https://doi.org/10.1177/0263276413486160>
- Steinberg, M. (2021). From automobile capitalism to platform capitalism: Toyotism as a prehistory of digital platforms. *Organization Studies*, 43(7), 1069–1090. <https://doi.org/10.1177/01708406211030681>
- Steinbock, J. (2022). Toward a political economy of synthetic data: A data-intensive capitalism that is not a surveillance capitalism. *New Media & Society*, 0(0), 1–17. <https://doi.org/10.1177/14614448221099217>
- Szymkowski, S. (2021, September 29). 2022 Cadillac Escalade drops hands-free Super Cruise system amid chip shortage. *CNET*. <https://www.cnet.com/roadshow/news/2022-cadillac-escalade-super-cruise-hands-free-chip-shortage>
- 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>
- Thrun, S., Montemerlo, M., Dahlkamp, H., Stavens, D., Aron, A., Diebel, J., Fong, P., Gale, J., Halpenny, M., Hoffman, G., Lau, K., Oakley, C., Palatucci, M., Pratt, V., Stang, P., Strohband, S., Dupont, C., Jendrossek, L.-E., Koelen, C., Markey, C., Rummel, C., van Niekerk, J., Jensen, E., Alessandrini P., Bradski, G., Davies, B., Ettinger, S., Kaehler, A., Nefian, A., & Mahoney, P. (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.
- Ting-Fang, C., & Li, L. (2022, August 5). 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>
- Urtasun, R. (2021, June 20). CVPR'21 WAD keynote – Raquel Urtasun, Waabi/University of Toronto. *Youtube*. <https://www.youtube.com/watch?v=PSZ2Px9PrHg>
- 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>
- WAD (2022). *Workshop on autonomous driving*. CVPR 2022. <https://cvpr2022.wad.vision/>
- WAD (2023). *Workshop on autonomous driving*. CVPR 2023. <https://cvpr2023.wad.vision/>
- 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>
- Waymo (2019, August 21). Waymo open dataset: Sharing our self-driving data for research. *Waypoint*. <https://blog.waymo.com/2019/08/waymo-open-dataset-sharing-our-self.html>
- Waymo (2020, March 19). Announcing Waymo's open dataset virtual challenges. *Waypoint*. <https://blog.waymo.com/2020/03/announcing-waymos-open-dataset-challenges.html>
- Wilken, R., & Thomas, J. (2019). Maps and the autonomous vehicle as a communication platform. *International Journal of Communication*, 13, 2703–2727.
- Zeng, W., Luo, W., Suo, S., Sadat, A., Yang, B., Casas, S., & Urtasun, R. (2021). End-to-end interpretable neural motion planner. *ArXiv*. <https://doi.org/10.48550/arXiv.2101.06679>