



The system of auto-mobility: computer vision and urban complexity—reflections on artificial intelligence at urban scale

Fabio Iapaolo¹

Received: 1 June 2021 / Accepted: 6 April 2022 / Published online: 8 May 2023
© The Author(s) 2023

Abstract

Focused on city-scale automation, and using self-driving cars (SDCs) as a case study, this article reflects on the role of AI—and in particular, computer vision systems used for mapping and navigation—as a catalyst for urban transformation. Urban research commonly presents AI and cities as having a one-way cause-and-effect relationship, giving undue weight to AI's impact on cities and overlooking the role of cities in shaping AI. Working at the intersection of data science and social research, this paper aims to counter this trend by exploring the reverse perspective: how do cities affect the development, and expose the present limits, of SDCs? The contribution of this paper is threefold. First, by comparing urban and nonurban environments and thoroughly examining the relationship between computer vision and city-specific sociality and form, it defines machine autonomy/automation as a function of the sociotechnical milieu in which an AI system operates. Second, and related, the paper problematizes the notion of SDCs as autonomous technologies and the role it plays in envisioning contending policy arrangements and technical solutions for achieving full driving automation. Finally, the article offers insight into a materialist and spatialized understanding of AI—namely, not as an abstract quality susceptible to replication within discrete machines, but rather as a distributed property emerging through embodied interactions among a multiplicity of agents (human, non-human, and technological) within/with their environments.

Keywords Self-driving cars · Computer vision · Urban AI · Auto-mobility · Machine learning · Autonomous technologies · Urban governance

1 Introduction

At the time of its first appearance on American public roads at the beginning of the twentieth century, the automobile was considered a luxury item reserved for the exclusive use and delight of a small number of wealthy enthusiasts and practitioners (Norton 2008). Since about the 1960s, the car has become ubiquitous (Dant 2004), establishing itself not just as an undisputed symbol of modernity (Lefebvre 1971), but also as one of the main catalysts for urban transformation—one capable of reshaping urban geography in its own image. Notoriously, for the automobile to become a mass medium of urban transportation (Dant and Martin 2001), profound cultural, regulatory, economic, and infrastructural changes were necessary over the past century (Featherstone

et al. 2005). Cities, in other words, had to be rearranged, socially and spatially, around cars (Norton 2008). For decades now, automobility has been considered “a common feature of everyday life itself, almost a background to the background” (Thrift 2004, pp. 45–46). However, the recent surge of interest in AI technologies has brought automobility back into the forefront of discussions about urban governance and planning. Such renewed attention is due primarily to the potential commercialization, in the near or distant future, of self-driving cars (SDCs)—namely, vehicles capable of automating all those functions that in traditional cars are governed by a human driver, including moment-to-moment decisions within high-stake, morally ambiguous situations. No longer perceived as futuristic objects from science fiction, SDCs are now a reality. Presently, positions about their possible introduction into urban roads are split. While some consider it inevitable (Claudel and Ratti 2015), others regard it as fantasy deemed to remain so, at least for a few decades (Casner et al. 2016; Janai et al. 2020). Either way, there is a general agreement on the great potential of

✉ Fabio Iapaolo
fiapaolo@brookes.ac.uk

¹ Faculty of Technology, Design and Environment, Oxford Brookes University, Oxford, UK

SDCs to radically reshape, as the motorcar did throughout the twentieth century, multiple aspects of urban life. How this will happen is currently subject to much debate and speculation among urban planners and engineers, architects, jurists, ethicists, and city governments.

For their advocates, including automakers, software providers, and institutional actors, SDCs have the potential to yield significant social gains, such as enhanced safety, accessibility, efficiency, and sustainability. Increased road safety, in particular, provides the primary justification for their political acceptability. According to oft-reported estimates, around 94% of all fatal crashes are due to human error. By removing human error from the driving equation, SDCs promise a drastic reduction in road accidents and fatalities. No longer requiring a person behind the wheel, SDCs are also believed to allow for more personal leisure, while expanding access to point-to-point mobility to previously underserved populations, such as the elderly and visually impaired (Woyke 2016). Another expected benefit is that, combined with emerging trends in car sharing, SDCs could ease traffic congestion with a projected 80% reduction in privately owned vehicles (Claudel and Ratti 2015). As a result, scarce urban assets like land and buildings used for parking could be reallocated for other purposes (Ratti and Biderman 2017).

Notwithstanding the potential benefits listed above, there are still many unanswered questions regarding their short- and medium-term effects, including how they will disrupt and reshape cities' political economy, public infrastructure, people's experience of space, and urban design itself. In this regard, it suffices to think of the potential cascading effects resulting from the marginalization or displacement of taxi drivers and other jobs currently requiring a human behind the wheel (Maughan 2019). Further eluding anticipation are the potential negative effects on broader political and socio-economic arrangements. The significant investments required for the city-wide deployment of SDCs will likely divert resources away from public transportation and other crucial policy domains, such as healthcare and education (Blyth et al. 2016). Widely debated are issues of liability and accountability in the event of accidents (Ganesh 2017). There is currently a regulatory gap concerning international safety and legal standards that makes it unclear who would be held legally responsible for any injuries or property damage caused by the vehicle, such as the owner, manufacturer, or code developers (Schellekens 2015). Additionally, the extensive use of software in SDCs raises concerns about the potential risks of hacking attacks (Maughan 2019).

About one century ago, the automobile (or 'horseless carriage' as it was also referred to at the time) promised to modernize a personal mobility system primarily reliant on animal-pulled vehicles, which car advocates considered ill-suited to modern society's needs (Norton 2008; Tarr

and Mcshane 2007). However, the extensive, far-reaching impacts, both intended and unintended (such as urban sprawl and air pollution), of the automobile were not fully understood until much later (Jacobs 1961). Currently, claims in favor of adopting SDCs from public and private actors, such as Waymo, Tesla, Uber, and others, reflect a vision for the future that seems both politically desirable and technologically inevitable. From their perspective, the transition from manual to fully automated vehicles is viewed primarily as a technological advancement publicly justified based on the expectation that SDCs will remedy the negative externalities associated with traditional cars.¹

This vision, however, is grounded upon the flawed assumption that the transition to full driving automation will be straightforward and without political tensions. SDCs are often marketed as a technology that will significantly ameliorate city life while maintaining existing socio-spatial systems substantially unchanged. However, this is unlikely to be the case. The mundane truth is that, as concisely yet powerfully stated by Stilgoe (2017b, p. 5), "[t]his plug-and-play story, in which the car is seen as able to get along with the world's complexities as they are, without making additional demands, is a lie". In other words, due to manifest incompatibilities between the technology's operating logics and current urban infrastructure and social behaviors, significant changes will be required before SDCs can be used in cities. Specifically, it is argued here, cities will be reshaped, socially and spatially, to accommodate the preemptive logics of machine vision systems used to map the vehicle's surroundings and predict future occurrences happening therein.

Focused on city-scale automation and using SDCs as a case study, this article reflects on the role of AI, and in particular computer vision—the main technical issue addressed and one of considerable practical policy and engineering import—as a catalyst for urban transformation.² So far, critiques of SDCs have mainly centered around ethical and legal ambiguities due to their autonomous operations and decision-making (Ganesh 2017); the contested governance

¹ The Society of Automotive Engineers International (2014) has proposed a general taxonomy for on-road motor vehicle classification, ranging from Level 0 (no automation) to Level 5 (full automation).

² The expression AI is used here in a broad sense. To date, no consensus has emerged about what AI is (not least because the term *intelligence* itself eludes univocal conceptualization). Depending on one's disciplinary perspective, analytical scale of interest (e.g., isolated algorithms or the broader sociotechnical systems of which they are but one element), and intellectual, political even, aims, AI can be understood in manifold ways. It is beyond my scope to provide a definition of AI. Yet, what needs to be emphasized is that AI is not just one single technology. Quite the opposite, the term is now commonly understood to be a general label for "a constellation of technologies" including, but not limited to, "machine learning, perception, reasoning, and natural language processing" (AI Institute 2016, p. 4).

of social innovation (Blyth et al. 2016; Stilgoe 2017a, 2017b; Tæiegh and Lim 2019; Marres 2020); the uneven distribution of power within complex human-technical assemblages (Bissell 2018; Ganesh 2020); the erosion of cognitive skills exteriorized into machines (Casner et al. 2020); and their broader societal impacts (Bissell et al. 2020). Only recently, scholars have started to examine the specific urban implications of SDCs (Duarte and Ratti 2018). Notwithstanding, apart from a few exceptions (Stilgoe 2017b; Bissell 2018; Iapaolo 2019; Bissell et al. 2020), urban research commonly presents SDCs and cities as having a one-way cause-and-effect relationship, putting excessive focus on AI's effect on cities and neglecting the impact of cities on AI. This paper challenges the prevailing techno-deterministic views by exploring the reverse perspective: how do cities affect the development—and expose the present limits—of SDCs?³

An SDC can be thought of as either a purely technical assemblage consisting of hardware and software that function together in executing the driving task (e.g., sensorimotor systems, computer vision, and machine learning algorithms), or as a sociotechnical assemblage embedded “within larger interlocking systems, rather than ... as [a] discrete entit[y]” (Bissell et al. 2020, p.10). An SDC, in other words, can be seen as part of what I define here, rephrasing what Urry (2005) wrote about traditional cars almost two decades ago, as a *system of auto-mobility*, which encompasses more than just vehicles but also includes physical and digital infrastructure, machine learning algorithms, training datasets, geolocation and mapping systems, three-dimensional cartographies, laws and codes of the road, roboethics, mobility cultures, governance structures, and new social habits and lifestyles.

The methodology of this article, which aims to bridge the gap between social research and computer science by combining advances from both fields, is based on a triangular emphasis that dialogically explores the political, spatial, and technical aspects of AI (see Hayles 1990).⁴ Following

Pasquinelli's (2019, p. 3, italic in the original) admonishment that existing research and debates on AI (see Ouchchy 2020) frequently “remai[n] at the level of speculation (*‘what if AI’*) and fai[l] at clarifying machine learning inner logic and intrinsic limits (*‘what is AI’*)”, this paper explores the underlying computational logics of SDCs in relation to city-specific sociality and form. Importantly, it does so by reintroducing materiality alongside technicality as key aspects for a better appreciation of the spatialized effects of AI.

Theoretical in scope and interdisciplinary in orientation, the contribution of this paper is threefold. First, by comparing urban and nonurban environments, it defines machine autonomy/automation as a function of the socio-technical milieu in which an AI system operates. Here, I endorse Kitchin's (2017, see also Hayles 2017) approach to researching algorithms, which posits that the focus should not be on algorithms alone but rather on the broader socio-technical assemblages they are part of, including “social practices, material properties, discourses, mathematical abstractions, and code” (Matzner 2019, p. 4). This perspective recognizes that automated actions and decisions are shaped by the intricate interplay between algorithms and the technical elements and subsystems they are networked with (e.g., sensorimotor subsystems). At the same time, they depend on contextual factors, arising through reciprocal interactions with the infrastructure of the built environment and the physical world of objects and people (Kitchin and Dodge 2011, see also Blanchette 2012). Drawing on the state-of-the-art in driving research, the article then delves into some of the key challenges currently hindering the introduction of SDCs into cities. By exposing how machine vision systems operate—and fail—within environments designed for the human senses, the paper problematizes the notion of SDCs as autonomous technologies and its role in envisioning contending policy arrangements and practical solutions for achieving full driving automation. Proposing the conceptual lens of *auto-mobility*, it is argued that a precondition for the city-wide deployment of SDCs will be the social and spatial reconfiguration of urban spaces to compensate for today's weaknesses of machine learning and vision—with requirements (social, political, spatial, and environmental) extending within and beyond the immediate urban surroundings. In the conclusive remarks, the article makes the argument for a materialist and spatialized understanding of AI—namely, not as an abstract quality susceptible to replication within discrete machines, but rather as a distributed property emerging through embodied interactions among a multiplicity of agents (human, non-human, and technological) within/with their sociotechnical environments.

³ As illustrated by Fuerth (2009, p. 20), the main problem with linearity as applied to the explanation of sociotechnical transformations is that it “distorts our notion of cause and effect. Under its influence, we tend to expect that for every problem there is a unique solution; and that proportionate changes of circumstances will produce proportionate changes of outputs. We believe that it is possible to disassemble (“unpack”) compound, conglomerate issues, without destroying their coherence”.

⁴ Technical knowledge on SDCs, and AI in general, has been gained through a systematic survey of refereed articles on the subject from engineer/computer science, college-level textbooks (e.g., Russel and Norvig, 2016), interviews with industry experts, as well as by visiting the research group: Artificial Intelligence and Media Philosophy (in German: *Künstliche Intelligenz und Medienphilosophie*, KIM), Karlsruhe University of Art and Design. For a technical survey on SDCs, see Yurtsever et al (2020). For a comprehensive overview on problems, datasets, and methods in computer vision for driving automation, see Janai et al (2020).

2 Artificial sensorium

In the last few years, SDCs have been trending in discussions of the future of mobility, “captur[ing] the popular imagination arguably more so than any other transportation technology over the past half century” (Bissell et al. 2020, p. 117). Nevertheless, or perhaps precisely because of that, an important aspect entirely overlooked in the current debates about self-driving technology is that automated transport systems are already fully operational in nonurban settings. Examples include automatic shuttles moving people between and within airport terminals or unmanned vehicles used for good transportation in industrial settings, such as factories, mines, and ports (see Chu et al. 2018, on the future of automated ports). These vehicles embody a degree of technological sophistication much less advanced than that SDCs do. Yet, unlike the latter, they are employed within standardized, controlled, and highly predictable environments, where the range of unexpected situations and human-related events the vehicle has to handle is very limited. An automatic train, for example, operates on a predetermined path with set stops, making its functions—and the task environment—predictable and thereby pre-programmable. Similarly, driverless vehicles utilized in industrial settings automate tasks in standardized and single-functional spaces with minimal-to-no interaction with other vehicles or people.

On the other hand, SDCs are designed to work in complex, non-predictable, and information-rich environments, which is exactly what urban roads are. On a busy road, an SDC must react within milliseconds to a wide range of situations that cannot be fully anticipated. City streets, due to their high density, diverse architecture, and variety of road users, add even more uncertainty to the environment. Traffic coordination can be seen as a classic instance of joint action (Chater et al. 2018), with mutual strangers negotiating intended actions and decisions according to shared rules. In dense city traffic, a car must simultaneously interact with a multitude of actors (e.g., pedestrians, bike runners, other vehicles, and animals), each acting independently in complicated ways. At any given moment, an SDC must, to mention but a few examples, continuously adapt to changing circumstances such as traffic speed; observe traffic rules; handle ambiguous situations (e.g., hand signals from traffic officers or construction workers); interpret unwritten conventions, and react to emergencies or rule-breaking behavior by other road participants (e.g., a pedestrian crossing the street in a non-designated section of the roadway).

Due to the high degree of uncertainty in their task environment and, consequently, the high unpredictability inherent to their operations, it has become standard practice to refer to SDCs as autonomous technologies. That is because

the way SDCs operate involves a wide *degree of indeterminacy*,⁵ here understood as the extent to which the vehicle’s behavior is continuously adaptive to contingent factors and situations, making it difficult to be governed through formal rules. SDCs operate in a way that is predictable at a high level (e.g., the vehicle will go from point A to point B), but not moment by moment, leading to the possibility of unexpected outcomes. To ensure responsiveness to external stimuli, SDCs use machine learning algorithms, sensorimotor technologies, and geo-referenced mapping and positioning systems. In concert, these technologies embody complex forms of cognition (Hayles 2017) that allow the vehicle to perceive and understand its surroundings and choose the best possible course of action from various alternatives. In this sense—and in this sense only—an SDC can be said to be operating *autonomously* from its designers and end-users.

Compared to traditional, logic-based forms of automation where outcomes are “already implicated in initial premises” (Parisi 2019, p. 3), SDCs have become paradigmatic of widespread conceptualizations of AI as discrete technologies or technical systems, the most advanced of which are said to be capable of “operat[ing] without the need for human intervention or supervision, mak[ing] decisions independently, and accommodat[ing] to changed circumstances” (Kaplan 2016, p. 147). Although it may be misleading, the common practice of extending autonomy beyond the human to include technology provides valuable insights. In this context, the term ‘autonomy’ encompasses both a political and technical meaning. Politically, it symbolizes the popular imagination, or rather misconception, of SDCs as “neoliberal, individualised agent[s] ... that can act independently and efficiently on the basis of guidelines and feedback” (Ganesh 2017, p. 7). Technically, it speaks for the advancement of sensing and cognitive abilities in current and developing AI technologies (see Lynch and Del Casino 2020; Hayles 2017).

Despite business claims (see Tennant and Stilgoe 2021) and popular ideas about technology’s autonomy (see Winner 1977), SDCs do not exist in a vacuum. They are quintessentially socio-spatial agents whose main goal is to transport passengers from one point to another. In the process, they navigate a world shared with many other entities (see Matern 2017): humans (pedestrians, cyclists), animals, objects (road signs, street furniture), and whatever elements they encounter during their journey. To safely perform their task and avoid collisions, SDCs must be able to sense and interpret their surroundings under any weather and lighting conditions. In a sense, SDCs instantiate what the French philosopher and urbanist Paul Virilio (1994, p. 59), writing in the Eighties, defined as *sightless vision*, in which “the capacity

⁵ The expression *degree of indeterminacy* is borrowed from the French philosopher Gilbert Simondon (1958/2017, p. 17).

to analyse the ambient environment and automatically interpret the meaning of events” is delegated to the dyad computer–camera. In driving automation, that which with anthropomorphic vocabulary is named machine perception, in reality, refers to complex statistical and probabilistic models used to identify the type and location of nearby objects and anticipate their behavior in the near future. Manovich (1996, p. 12) observes that “in the field of computer vision, “understanding a scene” implies two goals. First, it means the identification of various objects represented in an image. Second, it means reconstruction of three-dimensional space from image”.⁶ In SDCs, machine perception relies on four different types of sensors: cameras, radars, ultrasound sensors, and LiDAR scanners. Schematically, the overall decisional loop of an SDC follows a perception–decision–action sequence. During the perception stage, sensor readings are used to create a real-time spatiotemporal view of the vehicle itself in relation to its surroundings, including other vehicles, road signs, and street elements. Combined with prior knowledge of the road infrastructure and driving rules, this information is then used to calculate driving decisions that are executed by the vehicle's mechanical actuators.

It is important to note that exteroception (the perception of the road environment) is never 100% accurate. Just like human decision-making, decisions made by SDCs are always based on incomplete and imperfect information. Sources of uncertainty can come from either internal elements within the vehicle, such as faulty sensors or data processing mistakes, or external factors. To mitigate internal uncertainty, SDCs are equipped with multiple sensors of different types. The rationale is that fusing data from multiple sensor sources helps achieve more accurate results compared to relying on just one (Hall and Llinas 1997). However, in complex urban environments, the largest source of uncertainty is usually external and related to real-time coordination with other traffic participants and the road infrastructure.

Up until today, urban roads have been built around people and, with the rapid rise of automotive traffic, around motorists primarily. By “impos[ing] a strong social control over the most fundamental of human behaviors, whether to move or be still” (McShane 1999, p. 370), traffic signing systems (i.e., traffic lights, road signs, and painted pavement) are one crucial element of the road infrastructure. Used to convey helpful information for navigation—restrictions, prohibitions, and warnings—they dictate behavior and facilitate coordination among traffic participants. International

design principles for traffic signing systems are regulated under the Vienna Convention on Signs and Signals (1968). The convention categorizes road signs into seven classes, such as danger warning and priority signs, and specifies the color, size, and shape of each class. It also outlines the requirements for road markings, such as length, width, color, material, and message, as well as the colors and meanings of traffic lights.

Understandably, these standards have been defined to meet the demands of the human eye; that is, to minimize the time gap between the presentation of the stimulus and the driver's enactment of the appropriate response. Road signs, for instance, are strategically positioned for easy and unambiguous interpretation by drivers, particularly while driving at high speeds. By the same logic, road markings use light-colored retroreflective materials for maximum visibility during both day and night. During the day, they are visible due to their contrast against the dark pavement, while at night, retroreflective pigments bounce light back from the vehicle's headlights. Since the visual stimuli that mediate urban flows are designed for human and not machine vision, interpreting traffic signs and reacting to signaling are challenging tasks for SDCs, especially in busy city environments where significant computational resources are required to detect and track moving objects.

A partial solution to problem, one which, following Waymo (2016), is now common to most manufacturers, is the combination of real-time sensor data with centimeter-accurate maps that provide contextual knowledge on, for example, lane geometries, traffic rules, and location of traffic lights. This approach not only facilitates the vehicle's self-location in the absence of GPS, but also reduces the computational load otherwise needed for mapping permanent features of the built environment. Urban roads are, however, fast-changing environments (e.g., due to new infrastructure or regulations). Admittedly (see Waymo 2021), building, maintaining, and updating maps with such fine detail is a prohibitively costly—labor-intensive and time-consuming—activity. The process entails the use of human-operated vehicles equipped with LiDAR technology to capture 3D images of selected urban areas. These images must then undergo a meticulous manual review and validation before they can be utilized by SDCs.

The political significance of such an impressive cartographic project is easily recognizable (see Hardigree in 2019). Since SDCs are only capable of navigating pre-mapped and geo-fenced roads, these maps essentially define the geographical scope of the technology, limiting its adoption to those cities or urban areas deemed investment-worthy. Cartography, as highlighted by Denis Wood in his seminal book *The Power of Maps* (1992), is not just a scientific discipline but also a political endeavor aimed at securing mapmakers' control over territories. The high-resolution

⁶ The author (Manovich 1996, p. 12) further explains that “[a] robot, for instance, needs not only to recognize particular objects, but it has to reconstruct a representation of the surrounding environment to plan its movements”.

maps used for SDCs are not neutral either, but rather reflect the interests of their creators. This raises various questions, such as who has the authority to decide which areas should be mapped and made accessible, who sets the priorities, and how the process of making urban roads intelligible to SDCs will either reinforce existing or create new socio-economic divides within and between cities.

In addition to GPS signal disruption or loss due to tunnels or *urban canyons* (Cui and Ge 2003),⁷ two classic problems for computer vision in urban settings are road sign blockages caused by trees or other occasional occluding elements (Muio 2016) and reduced sensor performance in adverse lighting and weather conditions (Zang et al. 2019). The limitations of cameras and LiDAR technology in poor weather conditions, such as fog, heavy rain, or snowfall, are well documented (Dannheim et al. 2014). Additionally, minor issues that would not significantly affect human vision, such as faded or leaf-covered road markings, can cause problematic outcomes like off-course driving or unnecessary stops (Sage 2016; Flockett 2017; Barut 2018). Sporadic manifestations of latent flaws in computer vision include mistaking mannequins for pedestrians (Waymo 2021) or the moon for a yellow traffic light (Ramey 2021). Although some of these errors can be harmless, hilarious even, such as Tesla's autopilot confusing a Burger King sign for a stop sign (Lambert 2020), they can also pose a danger in safety-critical situations, as demonstrated by the fatal 2016 Tesla Model S accident (Yadron and Tynan 2016).

Just like any other technology, computer vision systems are susceptible to both deliberate and unintentional *adversarial attacks*. As shown in a study by Eykholt et al. (2018), even minor modifications to road signs, such as graffiti or stickers (Field 2017), can deceive computer vision systems. The researchers demonstrated that modest changes to a stop sign, like altering the background or adding graffiti, can result in the system mistaking it for a 45-mile-an-hour speed limit sign, causing the vehicle to speed up instead of stopping. In another recent experiment conducted within a controlled setting (Tencent Keen Security Lab 2019), stickers placed on the road prompted a Tesla Model S to switch lanes and drive abruptly into oncoming traffic. These studies highlight the vulnerability of current computer vision systems to even simple forms of deception. If such incidents were to occur in real-world scenarios, the consequences could be serious.

In densely populated urban areas, the interaction between SDCs and vulnerable road users, especially pedestrians and cyclists, can be particularly problematic. This is due to two main reasons. First, the unpredictable movements and changing directions of pedestrians and cyclists make it

difficult to predict their behavioral patterns. Second, people's appearance variability (e.g., physical attributes or clothing) negatively affects algorithms used for human detection and classification (Janai et al. 2020). These algorithms rely on supervised learning, where they are taught to recognize specific categories, such as differentiating between pedestrians and cyclists, through labeled examples. Indeed, before they can successfully identify anything, algorithms must first be instructed on what to see in the first place. However, in real-world scenarios, human unpredictability and variability limit their ability to apply their training effectively.

Also, in real-world situations, unexpected phenomena can emerge “that the training data simply did not include and could not have anticipated” (Gillespie 2014). The widespread use of masks during the COVID-19 pandemic highlights this issue. Commercial facial recognition systems like Apple Face ID, which had been trained on pre-pandemic images, suddenly had difficulties recognizing masked individuals (Simonite 2020). To address this problem, new training datasets were required, and in some cases, this was achieved by overlaying computer-generated masks on existing facial datasets (Ngan et al. 2020). Given that novel situations simply cannot be anticipated, it is likely that SDCs will face problems in the future comparable to those experienced by commercial software technology today—with even higher stakes involved.

These examples demonstrate that computer vision relies not only on algorithms and sensors, but also, and perhaps more significantly, on training datasets. Training datasets play a vital role in determining what SDCs can *see* and how they *behave*, even more so than the algorithms that govern the vehicle's behavior in real time (see Beer 2017 on the social power of algorithms). In driving automation, the quality of training datasets depends on various factors, such as the selection of useful taxonomies, categories, and subcategories; accuracy in data preparation and labeling; availability of training and test datasets reflecting all potential scenarios self-driving cars may face in real-world conditions, including *corner cases* (e.g., cars disregarding traffic signals)⁸; and the method of data collection, such as open-source datasets like ImageNet or ground truth data collected under actual traffic conditions (e.g., the KITTI dataset).⁹

⁸ The test dataset comprises data (e.g., images) that the learning algorithm has never encountered before. For this reason, test data must be distinct from data used for training and validation.

⁹ Unlike other publicly available domain-generic datasets like ImageNet, the KITTI dataset is specifically meant for computer vision applications in automated driving. The KITTI dataset, which is becoming an important benchmark in the field, is based on discrete images retrieved from video sequences recorded around the mid-size city of Karlsruhe, Germany, using a human-piloted car equipped with various sensors, including cameras and LiDAR. See <http://www.cvlibs.net/datasets/kitti/>

⁷ That is, city streets flanked by high-rise buildings on both sides.

These elements form the foundation of computer vision. They can be viewed as components in a machine learning assembly line (Pasquinelli and Joler 2020), the layered intricacies of which often conceal the fact that human decisions always already determine the extent and scope of machine vision. Far from merely defining the technological capabilities and limitations of computer vision, training datasets are, more evidently so than algorithms, imbued with values, assumptions, worldviews, or, to say it in one word, *politics*—in the sense of Winner (1980). Crawford and Paglen (2019) observe that "the automated interpretation of images is an inherently social and political project, rather than a purely technical one". Indeed, computer vision implies much more than just receiving and processing environmental cues; it also involves valuing and prioritizing certain things based on cultural norms and expectations. In this regard, if there is any important lesson to be learned from the famous "Trolley Problem" is that, contrary to its intended purpose (Awad et al. 2018, p. 59), there cannot be easily determined "global, socially acceptable principles" to guide engineers and programmers as values vary among cultures and circumstances (see Gold et al. 2014).

The central, and always open to negotiation (Amoore 2020), political challenge becomes determining what should be prioritized and made visible and what should be overlooked and invisibilised. In driving automation, nowhere is this more evident than in the initial specification of targets of interest (e.g., pedestrian, cyclist, or vehicle), though other factors also play a crucial role. For instance, SDCs have already been shown to have difficulty recognizing people with darker skin tones (Ganesh 2020, see Wilson et al. 2019). An unwanted consequence due to the uneven representation of different ethnic and social groups in the training data, this situation still illustrates the much-debated problem of social biases being amplified as they enter the realm of computation—in this case, by rendering minorities even less visible through the automation of perception tasks.

Until today, car manufacturers have primarily focused on upgrading their vehicles' technology to address the challenges of SDCs and mitigate the uncertainties that come with them. This has often involved equipping the cars with a more significant number and variety of sensors to ensure data redundancy but with the trade-off of greater computational costs and longer execution times. Nevertheless, as we will see later, it is becoming increasingly clear that vehicle-based solutions alone are not sufficient. The widespread integration of SDCs into cities will require not only technical advancements in areas such as machine learning and computer vision, but also a social and material restructuring of urban areas to create a supportive environment for SDCs.

3 The system of auto-mobility

State-of-the-art self-driving technology is not yet ready to handle urban complexity. Arguably, computer vision's unreliability in crowded city traffic is currently the main barrier to adopting SDCs in urban areas. Despite this, in the wake of what some have termed "testbed urbanism" (Halpern et al. 2013, see Kitchin 2014), in various cities around the world, there has been a push for street trials conducted by car manufacturers in partnership with local governments and other relevant stakeholders like software providers, universities, and private research centers. Famous examples include Waymo's partnership with the state of California (Subin and Wayland 2021) and Uber's partnership with Arizona, later temporarily suspended after a fatal pedestrian accident in 2018 (Bradshaw 2018).

Policies and legislation regarding SDCs vary from country to country and city to city, reflecting different, and sometimes conflicting, perspectives on driving automation. For example, in the United States, driving automation is primarily seen as the capability of individual vehicles to operate autonomously. This car-centered viewpoint holds that private technology companies, not government efforts, will drive the transition toward full driving automation by creating the required technology. In Europe, conversely, "the debate on connected cars still prevails over autonomous vehicles, feasibility of which is often questioned by EU officials under European driving conditions with complicated city centers" (FTI Consulting 2017, p. 2). The underlying idea is to develop cooperative/intelligent transport systems where vehicles are connected to each other and infrastructure through wireless technologies such as vehicle-to-vehicle and vehicle-to-everything communications (see Ionita 2017).¹⁰

Despite these variations, some trends are ubiquitous. In the absence of an international agreement on the regulation of SDCs, local governments are implementing regulations locally. To promote mutual benefits, so far, they have restricted their role to granting authorization, offering incentives, and potentially suspending on-road trials, while allowing major market players to take the lead in the research and development of SDCs. This "wait-and-see" approach (Grieman 2019) demonstrates confidence in the potential benefits of SDCs, such as improved safety and efficiency, but also assumes that the move toward full driving automation will not require significant public investment. On the other hand, for car manufacturers, gaining access to public roads can provide a significant competitive advantage and lead to lock-in effects, as the data gathered from real-life traffic conditions is proprietary.

¹⁰ One example being Barcelona 'superblock' initiative (Bausells 2017).

Due to a number of accidents involving semi-automated or fully automated vehicle prototypes, in recent years, car manufacturers have faced criticism for prioritizing their interests over public safety. More recently, however, there has been a shift in the public discourse, with increasing concerns being voiced by car manufacturers themselves and other key private players (see Stilgoe 2017a). Among these, one commonly held view is that the main factor presently hindering the widespread adoption of SDCs in cities is a lack of government funding for the necessary infrastructure. As a result, public scrutiny has shifted from questioning the capability of SDCs to safely navigate existing urban roads to problematizing cities for being ill-suited to the technology's requirements. The underlying idea is that, while there may still be some room for further advancement, self-driving technology has reached maturity. For Claudel and Ratti (2015), for instance, "[f]rom a technological point of view, driverless cars have arrived; the bigger task is for cities to integrate them". Likewise, Duval et al. (2019) argue that "[w]ithin cities, at least, a fully autonomous world awaits. Even though this world may be many years down the road, public officials should understand the changes ahead and consider the modifications needed to accommodate such systems". On the same lines, Oliver et al. (2018) maintain that "the key question we should be asking is not when will self-driving cars be ready for the roads, but rather which roads will be ready for self-driving cars". This almost unanimous call for greater involvement from public actors in the development and regulation of SDCs has two consequences. First, it redefines the political significance of those involved in governing and developing SDCs. Second, it lays the foundations for new practical solutions and policy arrangements for implementing SDCs in the short and long term.

As noted earlier, up until now, car manufacturers have primarily concentrated on enhancing the technology within their vehicles to cope with the challenges posed by urban areas. However, a more effective approach could be to alter the external environment to make it less complex for vehicles to navigate. This would entail transforming the urban landscape, including its design, material components, and social aspects, to overcome the limitations in computer vision technology. There are various ways to accomplish this, ranging from minor modifications to existing infrastructure to more ambitious plans for creating 'smart roads' that would optimize traffic flow through the use of wireless communication between vehicles and the road environment (Duvall et al. 2019). Ford, for instance, is testing a similar approach in Miami-Dade County, Florida. At critical intersections, the company has installed so-called 'smart nodes' equipped with cameras, radar, and LiDAR sensors. These nodes provide supplementary information to the data collected by the car's onboard sensors, such as other vehicles approaching a traffic light (Ford Motor Company 2021).

However, in a world where many municipalities struggle to maintain even ordinary roadways, the development of advanced smart infrastructure seems to be a distant goal for the time being. In the meantime, traditional vehicles and SDCs will likely coexist, yet with measures aimed at preventing or controlling their interactions. Lately, the idea of separating SDCs from other road users has gained traction, with various trials underway globally (Mullin 2020). This would involve creating dedicated lanes for SDCs with physical barriers separating them from human-operated vehicles (Krisher and Eggert 2020). Another solution could be limiting their use to slow-speed areas or structured environments such as urban highways (Hawkins 2018).

These approaches are more practical and economically viable compared to the extensive infrastructure changes needed in the long run. Nonetheless, the simplification of cities for SDCs will have an impact on both the built infrastructure and the social fabric of cities, potentially resulting in new forms of spatial segmentation and accompanying political conflicts. The history of the motor car highlights that urban roads have always been contested by many 'street rivals' competing for access to them (see Norton 2007). In the early 1900s, public streets were shared by various groups, including pedestrians, street vendors, children, and cyclists—all of whom had equal rights to the road. However, with the arrival of the motor car, cities witnessed the emergence of a new urban phenomenon, the car accident: "a new kind of mass death. Most of the dead were city people. Most of the car's urban victims were pedestrians, and most of the pedestrian victims were children and youths" (Norton 2008, p. 11). As cars, and their drivers, were initially seen, at best, as unruly intruders, city streets "had to be socially reconstructed as places where motorists unquestionably belonged" (Norton 2008, p. 1). Over time, the growing popularity of cars resulted in changes to city infrastructure, including the division of roads into specific areas for vehicles and others, more marginal, for pedestrians and cyclists. Along with these changes, new regulations were enacted to enforce these novel uses of the streets. In certain countries, including the United States, children were prohibited from playing on city streets, while the criminalization of 'jaywalking' restricted pedestrian freedom and rights. With the advent of SDCs, new forms of mobility and dwelling will emerge, but they will also require changes to current social practices and behaviors, including restrictions.

Urry (2004) pointed out that the widespread use of traditional cars as a personal mode of transportation was made possible by the concurrent growth of a globally expanding car system that gave them preference over other forms of mobility, such as walking, cycling, and rail. The author referred to this system as "automobility"—namely, "a self-organizing autopoietic, non-linear system that spreads world-wide, and includes cars, car-drivers, roads, petroleum

supplies and many novel objects, technologies and signs” (Urry 2004, pp. 26–27). Similarly, the widespread adoption of SDCs will require its own supporting system, which I refer to here as *autono-mobility*. This system encompasses vehicles and other elements, such as physical and digital infrastructure, hardware and software technology, machine learning algorithms, intelligent traffic control systems, training datasets, machine-readable maps, changes in liability and insurance, and the evolution of social norms and practices.

It should be remarked that the scope of the said system extends within and beyond the urban scale. The ambitious project to introduce SDCs into urban roads is akin to a ‘planetary experiment’ (see Halpern 2019), with a wide range of potential implications not just for cities but political, economic, and spatial systems at multiple interrelated scales. As a result, it is crucial to conduct a comprehensive evaluation of SDCs that considers aspects often overlooked in discussions focused solely on cities.

At present, conversations about urban automation primarily center on the potential for AI to replace humans in driving and other domains. However, this narrow perspective disregards the disproportionate amounts of planetary resources required for that to occur. In driving automation, one such resource is the screen workers responsible for annotating images and videos used to train the computer vision systems of SDCs. These workers play a critical role in creating the data infrastructure that supports the computer vision systems in SDCs. Still, this infrastructure also comes with high environmental costs (see Sudhakar et al. 2023). Thus, while it is crucial to interrogate and possibly anticipate the impact of SDCs on cities, it is just as important to recognize the existing demands they already place on a global network of resources, labor, and data (Crawford and Joler 2018).

4 Conclusion: on the spatiality and materiality of urban AI

This article has thoroughly examined the interplay between computer vision and city-specific sociality and spatiality. By taking a broader view of the sociotechnical system of autono-mobility, rather than just focusing on SDCs as the central unit of analysis, the article has made the case that urban transformation should not be solely viewed as a result of the integration of SDCs into cities, but rather as a necessary prerequisite for their successful deployment. Specifically, cities will need to undergo significant changes, socially and spatially, aimed at accommodating the requirements and limitations of today’s computer vision systems. Additionally, the study has advocated for a more comprehensive examination of the far-reaching, more-than-urban consequences of autono-mobility, including the planetary—social and

environmental—costs implicated in the ongoing process of bringing fully automated vehicles onto city streets.

Yet, another important aspect needs to be emphasized: SDCs are nowhere near as autonomous as they are commonly portrayed. This is not to undermine their advanced cognitive abilities, as demonstrated by their capability to understand contextual information and adjust to changing surroundings (Hayles 2017). However, if applied uncritically to AI and SDCs specifically, the concept of autonomy is complicated for three main reasons.

First, in the near term, the most feasible approach to integrating SDCs into urban roads is limiting their functions by restricting their usage to more structured environments such as dedicated lanes. This approach would make the vehicle’s behavior more predictable and, in a sense, less autonomous, similar to that of automated shuttles in airport complexes. Nonetheless, the idea of lowering the world’s complexities to increase technological autonomy is, at best, paradoxical. The same holds for long-term strategies. Currently, there is no consensus on the best way to manage the transition to full driving automation, except that it will require increased connectivity and interdependence between road actors to an unprecedented degree. In other words, enhancing road safety and improving traffic flow will be accomplished less through greater autonomy of individual vehicles and more through improving the interconnectivity among cars and between cars and road infrastructure, resulting in environments with collective cognitive capabilities that surpass those of individual actors.

Second, autonomy and intelligence are frequently mistaken as synonymous in discussions of AI, both among the general public and in academic circles, as if these traits are inherent in individuated technologies. This widespread perception of AI closely aligns with the liberal conception of humans as autonomous individuals possessing free will and innate intellectual faculties. However, this view disregards the extent to which intelligence is relational, embedded, situated, and infrastructural (Bruder 2019). Within an AI system, intelligence (and decision-making) can rarely, if ever, be attributed to one single entity. Instead, intelligent behavior is the outcome of complex interactions between multiple components and subsystems, each with its own material properties and limitations that influence the overall functioning of AI. At the same time, the situated decisions made by AI technologies, especially embodied systems like SDCs, are not made in isolation but shaped through mutual interactions with various human and non-human actors within a shared environment. Furthermore, as my comparative examination of urban and nonurban environments demonstrates, space is not just an incidental milieu but an integral aspect of AI that continuously shapes the affordances and limitations of the technologies it encompasses and interacts with.

Finally, it is essential to remember that AI is a human endeavor, encompassing all aspects from data collection to preparing training datasets and selecting the appropriate deployment environment. When humans disappear from the driving seat, it becomes ever more crucial to understand where else they are involved, including the relevant sites and scales where human decision-making continues to govern technology and affect its outcomes.

Acknowledgements A part of this paper appeared in Italian, in a significantly different version, in Iapaolo F (2019) *Città post-antropocentriche. Visione Artificiale e Complessità Urbana*. *Atti & Rassegna Tecnica* 73(2):177–183. I am sincerely grateful to the editors and two anonymous reviewers for their constructive commentaries.

Funding This paper was partially supported by Oxford Brookes University and by the Inter-University Department of Regional and Urban Studies and Planning (DIST), jointly activated by the Polytechnic of Turin and the University of Turin.

Data availability Not applicable.

Code availability Not applicable.

Declarations

Conflict of interest The author acknowledges that this work was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- AI Now Institute (2016) AI Now 2016 Report. https://ainowinstitute.org/AI_Now_2016_Report.pdf. Accessed 23 Mar 2018
- Amoore L (2020) *Cloud ethics: algorithms and the attributes of ourselves and others*. Duke University Press, Durham
- Awad E, Dsouza S, Kim R, Schulz J, Henrich J, Shariff A, Bonnefon JF, Rahwan I (2018) The moral machine experiment. *Nature* 563:59–64. <https://doi.org/10.1038/s41586-018-0637-6>
- Barut J (2018) This War-Proven Radar Tech Lets Self-Driving Cars See the Road When You Can't. Raod and Track. <https://www.roadandtrack.com/new-cars/car-technology/a22774804/waves-ense-ground-penetrating-radar-self-driving-autonomous-cars/>. Accessed 23 June 2019
- Bausells, M (2016) Superblocks to the rescue: Barcelona's plan to give streets back to residents. *Guardian*. <https://www.theguardian.com/cities/2016/may/17/superblocks-rescue-barcelona-spain-plan-give-streets-back-residents>. Accessed 4 Sept 2021
- Beer D (2017) The social power of algorithms. *Inf Commun Soc* 20(1):1–13. <https://doi.org/10.1080/1369118X.2016.1216147>
- Bissell D (2018) Automation interrupted: How autonomous vehicle accidents transform the material politics of automation. *Polit Geogr* 65:57–66
- Bissell D, Birtchnell T, Elliott A, Hsu EL (2020) Autonomous automobiles: The social impacts of driverless vehicles. *Curr Sociol* 68(1):116–134. <https://doi.org/10.1177/0011392118816743>
- Blanchette JF (2012) Computing as if infrastructure mattered. *Commun ACM* 55(10):32–34. <https://doi.org/10.1145/2347736.2347748>
- Blyth PL, Mladenovic NM, Nardi BA, Ekbia HR, Su NM (2016) Expanding the design horizon for self-driving vehicles: distributing benefits and burdens. *IEEE Technol Soc Mag* 35(3):44–49. <https://doi.org/10.1109/MTS.2016.2593199>
- Bradshaw T (2018) Self-driving cars under scrutiny after Uber pedestrian death. *Financial Times*. <https://www.ft.com/content/6d328aa8-2be5-11e8-9b4b-bc4b9f08f381>. Accessed 4 Oct 2021
- Bruder J (2019) *Cognitive code. Post-anthropocentric Intelligence and the Infrastructural Brain*. McGill-Queen's University Press, Kingston
- Casner M, Hutchins EL, Don N (2016) The challenges of partially automated driving. *Commun ACM* 59(5):70–77. <https://doi.org/10.1145/2830565>
- Chater N, Misyak J, Watson D, Griffiths N, Mouzakitis A (2018) Negotiating the traffic: Can cognitive science help make autonomous vehicles a reality? *Trends Cogn Sci* 22(2):93–95. <https://doi.org/10.1016/j.tics.2017.11.008>
- Chu F, Gailus S, Liu L, Ni L (2018) The future of automated ports. McKinsey. <https://www.mckinsey.com/industries/travel-logistics-and-infrastructure/our-insights/the-future-of-automated-ports>. Accessed 3 Oct 2020
- Claudel M, Ratti C (2015) Full speed ahead: How the driverless car could transform cities. McKinsey. <http://www.mckinsey.com/business-functions/sustainability-and-resource-productivity/our-insights/full-speed-ahead-how-the-driverless-car-could-transform-cities>. Accessed 23 Mar 2018
- Crawford K, Joler V (2018) Anatomy of an AI System: The Amazon Echo as an anatomical map of human labor, data and planetary resources. <https://anatomyof.ai/>. Accessed 22 May 2021
- Crawford K, Paglen T (2019) Excavating AI: The Politics of Images in Machine Learning Training Sets. <https://excavating.ai/>. Accessed 9 Sept 2020
- Cui Y, Ge SS (2003) Autonomous vehicle positioning with GPS in urban canyon environments. *IEEE Trans Robotics Autom* 19(1):15–25. <https://doi.org/10.1109/TRA.2002.807557>
- Dannheim C, Icking C, Mader M, Sallis PJ (2014). Weather Detection in Vehicles by Means of Camera and LIDAR Systems. 2014 Sixth International Conference on Computational Intelligence, Communication Systems and Networks, pp 186–191
- Dant T (2004) The Driver-car. *Theory Cult Soc* 21(4/5):61–79. <https://doi.org/10.1177/0263276404046061>
- Dant T, Martin P (2001) By car: carrying modern society. In: Warde A, Grunow J (eds) *Ordinary consumption*. Routledge, London, pp 143–157
- Duarte F, Ratti C (2018) The impact of autonomous vehicles on cities: a review. *J Urban Technol* 25(4):3–18. <https://doi.org/10.1080/10630732.2018.1493883>
- Duvall T, Hannon E, Katseff J, Safran B, Wallace T (2019) A new look at autonomous-vehicle infrastructure. McKinsey. <https://www.mckinsey.com/industries/travel-logistics-and-infrastructure/our-insights/a-new-look-at-autonomous-vehicle-infrastructure>. Accessed 30 Jan 2020
- Eykholt K, Evtimov I, Fernandes E, Li B, Rahmati A, Xiao C, Prakash A, Kohno T, Song DX (2018). Robust Physical-World Attacks on Deep Learning Visual Classification. arxiv preprint. <https://arxiv.org/abs/1707.08945v5>. Accessed 17 June 2020

- Featherstone M, Thrift N, Urry J (eds) (2005). Sage, London
- Field M (2017) Graffiti on stop signs could trick driverless cars into driving dangerously. Telegraph. <https://www.telegraph.co.uk/technology/2017/08/07/graffiti-road-signs-could-trick-driverless-cars-driving-dangerously/>. Accessed 19 Dec 2020
- Flockett A (2017) Marking roads to make them safer for self-driving cars. Electronic Specifier. <https://www.electronicspecifier.com/products/artificial-intelligence/marking-roads-to-make-them-safer-for-self-driving-cars>. Accessed 14 June 2020
- Ford Motor Company (2021) Exploring How Smart Infrastructure Can Help Ford Build a Great Self-Driving Service for Miamians. Medium. <https://medium.com/self-driven/exploring-how-smart-infrastructure-can-help-ford-build-a-great-self-driving-service-for-miamians-d1f6fccd0c14>. Accessed 3 Sept 2021
- FTI Consulting (2017) A Global Race for Autonomous Vehicles: Views from the United States, Europe and Asia. <https://euagenda.eu/upload/publications/untitled-91781-ea.pdf>. Accessed 27 Sept 2021
- Fuerth LS (2009) Foresight and anticipatory governance. *Foresight* 11(4):14–32. <https://doi.org/10.1108/14636680910982412>
- Ganesh MI (2017) Entanglement: machine learning and human ethics in driverless car crashes. *A Peer-Reviewed Journal about* 6(1):77–87
- Ganesh MI (2020) The ironies of autonomy. *Humanit Soc Sci Commun*. <https://doi.org/10.1057/s41599-020-00646-0>
- Gillespie, T (2014) Algorithm [draft] [#digitalkeywords]. *Culture Digitally*. <https://culturedigitally.org/2014/06/algorithm-draft-digitalkeyword/>. Accessed 23 Nov 2019.
- Gold N, Colman AM, Pulford BD (2014) Cultural differences in responses to real-life and hypothetical trolley problems. *Judgm Decis Mak* 9(1):65–76
- Grieman K (2019) Hard drive crash: An examination of liability for self-driving vehicles. *JIPITEC – Journal of Intellectual Property, Information Technology and E-Commerce Law*. <https://www.jipitec.eu/issues/jipitec-9-3-2018/4806>. Accessed 24 May 2020
- Hall DL, Llinas J (1997) An introduction to multisensor data fusion. *Proc IEEE* 85(1):6–23. <https://doi.org/10.1109/5.554205>
- Halpern O, LeCavalier J, Calvillo N, Pietsch W (2013) Test-bed urbanism. *Publ Cult* 25(2):272–306
- Hardigree M (2019) Geofences: the invisible walls surrounding autonomous cars. *Apex one*. <https://www.apex.one/articles/geofences-the-invisible-walls-surrounding-autonomous-cars/>. Accessed 23 Aug 2021
- Hawkins AJ (2018) Senior citizens will lead the self-driving revolution. *The Verge*. <https://www.theverge.com/2018/1/10/16874410/voynage-self-driving-cars-villages-florida-retirement-communities>. Accessed 23 June 2021
- Hayles NK (1990) *Chaos bound: orderly disorder in contemporary literature and science*. Cornell University Press, Ithaca
- Hayles NK (2017) *Untought: The power of the cognitive nonconscious*. University of Chicago Press, Chicago and London
- Iapaolo F (2019) Città post-antropocentriche. *Visione Artificiale e Complessità Urbana. Atti & Rassegna Tecnica* 73(2):177–183
- Ionita S (2017) Autonomous vehicles: from paradigms to technology. *IOP Conference Series: Materials Science and Engineering* 252. <https://doi.org/10.1088/1757-899X/252/1/012098>
- Jacobs J (1961) *Life and death of American cities*. Random House, New York
- Janai J, Güney F, Behl A, Geiger A (2020) Computer vision for autonomous vehicles: Problems, datasets and state of the art. *Foundations and Trends® in Computer Graphics and Vision* 12(1–3):1–308. <https://doi.org/10.1561/06000000079>
- Kaplan J (2016) *Artificial Intelligence: What everyone needs to know*. Oxford University Press, New York
- Kitchin R (2014) The real-time city? Big data and smart urbanism. *GeoJournal* 79:1–14. <https://doi.org/10.1007/s10708-013-9516-8>
- Kitchin R (2017) Thinking critically about and researching algorithms. *Inf Commun Soc* 20(1):14–29. <https://doi.org/10.1080/1369118X.2016.1154087>
- Kitchin R, Dodge M (2011) *Code/space: software and everyday life*. MIT Press, Cambridge
- Krisher T, Eggert T (2020) Michigan plans dedicated road lanes for autonomous vehicles. *Abc News*. <https://abcnews.go.com/Technology/wireStory/michigan-plans-dedicated-road-lanes-autonomous-vehicles-72352758>. Accessed 24 Aug 2021
- Lambert F (2020) Tesla Autopilot confuses Burger King for stop signs, they turn into ad campaign. *Electrek*. <https://electrek.co/2020/06/25/tesla-autopilot-confuses-burger-king-stop-signs-ad-campaign/>
- Lefebvre H (1971) *Everyday life in the modern world*. Penguin, London
- Lynch CR, Del Casino VJ (2020) Smart spaces, information processing, and the question of intelligence. *Ann Am Assoc Geogr* 110(2):382–390. <https://doi.org/10.1080/24694452.2019.1617103>
- Manovich L (1996) Automation of sight: From Photography to computer vision. In: Druckrey T (ed) *Electronic culture: Technology and visual representation*. Aperture, New York, pp 229–236
- Marres N (2020) What if nothing happens? Street trials of intelligent cars as experiments in participation. In: S Maassen, S Dickel, CH Schneider (eds) *TechnoScience in Society, Sociology of Knowledge Yearbook*. Springer/Kluwer, Nijmegen, pp 1–10
- Mattern S (2017) Mapping’s intelligent agents. *Places Journal*. <https://placesjournal.org/article/mappings-intelligent-agents/>. Accessed 27 Jan 2020
- Matzner T (2019) The Human Is Dead – Long Live the Algorithm! *Human-Algorithmic Ensembles and Liberal Subjectivity. Theory, Culture & Society* 0(0):1–22. <https://doi.org/10.1177/0263276418818877>
- Maughan T (2019) No one’s driving: Autonomous vehicles will reshape cities, but is anyone taking control of how? In: Young L (ed) *Machine landscapes: Architectures of the Post-Anthropocene*. John Wiley & Sons, Oxford, pp 92–99
- McShane C (1999) The origins and globalization of traffic control signals. *J Urban Hist* 25(3):379–404. <https://doi.org/10.1177/009614429902500304>
- Mullin L (2020) China’s cities will soon be crawling with self-driving robotaxis. *Fast Company*. <https://www.fastcompany.com/90553995/chinas-cities-will-soon-be-crawling-with-self-driving-robotaxis>. Accessed 14 Sept 2021
- Muoio D (2016) Uber’s self-driving cars are impressive—but there’s still a lot they can’t do. *Business Insider*. <https://www.businessinsider.com/uber-driverless-car-problems-2016-9?r=US&IR=T>. Accessed 16 May 2020
- Ng A, Lin Y (2016) Self-Driving Cars Won’t Work Until We Change Our Roads—And Attitudes. *Wired*. <https://www.wired.com/2016/03/self-driving-cars-wont-work-change-roads-attitudes/>. Accessed 7 Sept 2020
- Ngan M, Grother P, Hanaoka K (2020) Ongoing Face Recognition Vendor Test (FRVT) Part 6A: Face recognition accuracy with masks using pre- COVID-19 algorithms. <https://doi.org/10.6028/NIST.IR.8311>. Accessed 10 Oct 2021
- Norton PD (2007) Street rivals: jaywalking and the invention of the motor age street. *Technol Cult* 48(2):331–359. <https://doi.org/10.1353/tech.2007.0085>
- Norton PD (2008) *Fighting traffic: the dawn of the motor age in the American city*. MIT Press, Cambridge and London
- Oliver N, Potočník K, Calvard T (2018) To make self-driving cars safe, we also need better roads and infrastructure. *Harvard Business Review*. <https://hbr.org/2018/08/to-make-self-driving-cars-safe-we-also-need-better-roads-and-infrastructure>. Accessed 22 May 2019

- Halpern Orit (2019) The Planetary Test. In: ZMK Zeitschrift für Medien- und Kulturforschung. Ontography, 10(1):13–22. <https://doi.org/10.25969/mediarep/18717>
- Ouchchy L, Coin A, Dubljević V (2020) AI in the headlines: the portrayal of the ethical issues of artificial intelligence in the media. *AI & Society* 35(4):927–936. <https://doi.org/10.1007/s00146-020-00965-5>
- Parisi L (2019) Critical computation: digital automata and general artificial thinking. *Theory Cult Soc* 36(2):89–121. <https://doi.org/10.1177/0263276418818889>
- Pasquinelli M, Joler V (2020) The Noosope manifested: AI as instrument of knowledge extractivism. *AI & Soc.* <https://doi.org/10.1007/s00146-020-01097-6>
- Pasquinelli M (2019) How a machine learns and fails—A grammar of error for Artificial Intelligence. *Spheres*. <https://spheres-journal.org/contribution/how-a-machine-learns-and-fails-a-grammar-of-error-for-artificial-intelligence/>. Accessed 16 Apr 2020
- Ramey J (2021) Tesla FSD Mistakes Moon for Yellow Traffic Light. *Autoweek*. <https://www.autoweek.com/news/green-cars/a37114603/tesla-fsd-mistakes-moon-for-traffic-light/>. Accessed 23 Sept 2021
- Ratti C, Biderman A (2017) From parking lot to paradise. *Sci Am* 317(1):54–59. <https://doi.org/10.1038/scientificamerican0717-54>
- Russell SJ, Norvig P (eds) (2016) 3rd edn. Prentice Hall, Upper Saddle River
- Sage A (2016) Where's the lane? Self-driving cars confused by shabby U.S. roadways. *Reuters*. <https://www.reuters.com/article/us-autos-autonomous-infrastructure-insig-idUSKCN0WX131>. Accessed 13 May 2020
- Schellekens M (2015) Self-driving cars and the chilling effect of liability law. *Comput Law Sec Rev* 31(4):506–517. <https://doi.org/10.1016/j.clsr.2015.05.012>
- Schmidt FA (2019) Crowdsourced production of AI Training Data: How human workers teach self-driving cars how to see, Working Paper Forschungsförderung, No. 155, Hans-Böckler-Stiftung, Düsseldorf, <https://nbn-resolving.de/urn:nbn:de:101:1-2019102414513364122668>. Accessed 9 May 2020
- Simondon G (2017) On the Mode of Existence of Technical Objects (C Malaspina and J Rogove, Trans.) University of Minnesota Press, Minneapolis (Original work published 1958)
- Simonite T (2020) How Well Can Algorithms Recognize Your Masked Face?. *Wired*. <https://www.wired.com/story/algorithms-recognize-masked-face/>. Accessed 24 Feb 2021
- Society of Automotive Engineers International (2014) Taxonomy and definitions for terms related to on-road motor vehicle automated driving systems J3016_201401. https://www.sae.org/standards/content/j3016_201401/preview/. Accessed 25 May 2019
- Stilgoe J (2017a) Machine learning, social learning and the governance of self-driving cars. *Soc Stud Sci* 18(4):725–748. <https://doi.org/10.1177/0306312717741687>
- Stilgoe J (2017b) Seeing like a Tesla: How can we anticipate self-driving worlds?. *Glocalism* 3:1–20. <https://doi.org/10.12893/gjcp.2017.3.2>
- Subin S, Wayland M (2021) Alphabet's Waymo and GM's Cruise get California DMV approval to run commercial autonomous car services. *CNBC*. <https://www.cnbc.com/2021/09/30/waymo-and-cruise-get-california-dmv-approval-to-run-driverless-cars.html>. Accessed 7 Oct 2021
- Sudhakar S, Sze V, Karaman S (2023) Data centers on wheels: emissions from computing onboard autonomous vehicles. *IEEE Micro* 43(1):29–39. <https://doi.org/10.1109/MM.2022.3219803>
- Taeihagh A, Lim HSM (2019) Governing autonomous vehicles: emerging responses for safety, liability, privacy, cybersecurity, and industry risks. *Transp Rev* 39(1):103–128. <https://doi.org/10.1080/01441647.2018.1494640>
- Tarr JA, Mcshane C (2007) *The horse in the city: living machines in the nineteenth century*. Johns Hopkins University Press, Baltimore
- Tencent Keen Security Lab (2019) Experimental security research of tesla autopilot. https://keenlab.tencent.com/en/whitepapers/Experimental_Security_Research_of_Tesla_Autopilot.pdf. Accessed 13 Sept 2021
- Tennant C, Stilgoe J (2021) The attachments of 'autonomous' vehicles. *Soc Stud Sci*. <https://doi.org/10.1177/03063127211038752>
- Thrift N (2004) Driving in the city. *Theory Cult Soc* 21(4/5):41–59. <https://doi.org/10.1177/0263276404046060>
- Urry J (2005) The 'system' of automobility. In: Featherstone M, Thrift N, Urry J (eds) *Automobilities*. Sage, London, pp 25–39
- Vienna Convention on Road Signs and Signals (1968) https://unece.org/DAM/trans/conventn/Conv_road_signs_2006v_EN.pdf. Accessed 28 May 2021
- Virilio P (1994) *The vision machine*. Indiana University Press, Bloomington and Indianapolis
- Waymo (2016) Building maps for a self-driving car. *Medium*. <https://medium.com/waymo/building-maps-for-a-self-driving-car-723b4d9cd3f4>. Accessed 23 Sept 2021
- Waymo (2021) The Waymo Driver Handbook: How our highly-detailed maps help unlock new locations for autonomous driving. *Blog Waymo*. <https://blog.waymo.com/2020/09/the-waymo-driver-handbook-mapping.html>. Accessed 23 Sept 2021
- Webb A (2017) Distributed cognition: Assessing the structure of urban scale Artificial Intelligence. *International Robotics & Automation Journal* 2(5):187–194. <https://doi.org/10.15406/iratj.2017.02.00034>
- Wilson B, Hoffman J, Morgenstern J (2019) Predictive inequity in object detection. *arXiv:1902.11097v1 [cs.CV]*. Accessed 25 May 2020
- Winner L (1977) *Autonomous technology: technics-out-of-control as a theme in political thought*. MIT Press, Cambridge
- Winner L (1980) Do Artifacts Have Politics? *Daedalus* 109(1):21–36
- Wood D (1992) *The Power of Maps*. Guilford, New York
- Woyke E (2016) The blind community has high hopes for self-driving cars. *MIT Technology Review*. <https://www.technologyreview.com/s/602555/the-blind-community-has-high-hopes-for-self-driving-cars/>. Accessed 28 May 2018
- Yadron and Tynan (2016) Tesla driver dies in first fatal crash while using autopilot mode. <https://www.theguardian.com/technology/2016/jun/30/tesla-autopilot-death-self-driving-car-elon-musk>. Accessed 23 June 2020
- Yurtsever E, Lambert J, Carballo A, Takeda K (2020) A Survey of autonomous driving: Common practices and emerging technologies. *IEEE Access* 8:58443–58469. <https://doi.org/10.1109/ACCESS.2020.2983149>
- Zang S, Ding M, Smith DK, Tyler P, Rakotoarivelo T, Kâafar MA (2019) The impact of adverse weather conditions on autonomous vehicles: how rain, snow, fog, and hail affect the performance of a self-driving car. *IEEE Veh Technol Mag* 14:103–111. <https://doi.org/10.1109/MVT.2019.2892497>

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.