On the future of transportation in an era of automated and autonomous vehicles

P. A. Hancock[a,b,1], Illah Nourbakhsh[c], and Jack Stewart[d]

[a]Department of Psychology, University of Central Florida, Orlando, FL 32816; [b]Institute for Simulation and Training, University of Central Florida, Orlando, FL 32826; [c]The Robotics Institute, Carnegie Mellon University, Pittsburgh, PA 15213; and [d]Wired Magazine, San Francisco, CA 94107

Edited by Baruch Fischhoff, Carnegie Mellon University, Pittsburgh, PA, and approved November 1, 2018 (received for review April 3, 2018)

Automated vehicles (AVs) already navigate US highways and those of many other nations around the world. Current questions about AVs do not now revolve around whether such technologies should or should not be implemented; they are already with us. Rather, such questions are more and more focused on how such technologies will impact evolving transportation systems, our social world, and the individuals who live within it and whether such systems ought to be fully automated or remain under some form of direct human control. More importantly, how will mobility itself change as these independent operational vehicles first share and then dominate our roadways? How will the public be kept apprised of their evolving capacities, and what will be the impact of science and the communication of scientific advances across the varying forms of social media on these developments? We look here to address these issues and to provide some suggestions for the problems that are currently emerging.

automated vehicles | future transportation infrastructure | trust in autonomy | public reactions | individuated technology

Automated vehicles (AVs) are coming to our roadways. They are beginning to pose problems and issues that many of the public have not yet encountered or even witnessed. The present work addresses how the capacities and impacts of growing degrees of ground automation can be communicated to and understood by the general public. To accomplish this, we have sought to integrate our individual perspectives on the issue of AVs which feature, first, a science communication viewpoint that focuses specifically on how the traveling public can anticipate, understand, and appreciate the effects of such innovation. Our second narrative strand features a human-centered approach to the on-coming penetration of AVs, looking to understand precisely how these diverse forms of full and semiautomation will be experienced by human drivers. Finally, we conclude with an analysis of the technical challenges guiding the possible features of this wave of automation and prospective autonomy in future transportation. All three levels—public communication, human–machine interaction, and technical feasibility—coact to sculpt the coming forms of transportation. The resulting system promises to be strikingly different from its traditional and contemporary form, which has come to be accepted as the status quo for almost a century. Shared discourse, including public communications pertaining to this disruptive evolution, is critical to our collective understanding of the future we may be able to create.

Current Communications About Self-Driving Cars

As of March 2018, 52 companies possessed permits to test autonomous vehicles on the roads of the State of California alone (1, 2). Self-driving vehicles represent a fast-paced field of modern technology, as companies compete for dominance in this important field of emerging transportation capacity. Nevertheless, relatively few members of the traveling public have yet experienced trips in an autonomous vehicle. This personal inexperience can make it difficult for the general populace to judge the potential utility, for good or bad, of such vehicles. The advent of the driverless car is usually portrayed as both labor saving and accident reducing. However, the societal impact of these mobile robots will certainly be more extensive than a simple change in the journey between the immediate origin and the destination. For example, in coming years it may not be necessary for individuals to own a car, especially when they can summon one from a circulating fleet using a simple portal such as a smartphone application and being fully confident that it will arrive within minutes or even seconds. This sea change in vehicle usage will have many knock-on effects. Some studies have suggested that up to 30% or more of traffic circling downtown streets is actually searching for parking (3). The search could become unnecessary when the vehicle is driving itself to pick up its next user, as some projections concerning Uber usage seem to suggest. Such functionality could free up curb space, which is becoming increasingly more important for safe pick-ups and drop-offs in already congested locations (4). These technological changes may then foreshadow a repurposing of parking structures or parking spaces within buildings to accommodate new housing, offices, or retail uses. Of course, parking concerns are not by any means the only dimension of change.

The radical changes promised by AVs will have profound and extended effects on the general public. Some of these are changes we can readily anticipate; others are much less predictable. On what basis will individual members of the public judge the value of such technical innovations (5)? One prominent issue in such a discussion is what people understand an AV to be. It may well be that the general public views such vehicles as not requiring any driver input whatsoever. However, this perception fails to capture many of the major differences between proposed AVs and the present, semiautomated on-road vehicles. The latter provide various forms of driver assistance to help the driver who remains in ultimate control. Fully autonomous vehicles are designed to drive themselves. These differing forms of advancing vehicles have been categorized in a hierarchy which compares driver control versus vehicle control. The hierarchy is described in the Society of Automotive Engineers (SAE) levels of control (6, 7). Although we do not specifically discuss each of these levels here, it is vital to note that many public assumptions about advanced vehicle capabilities may be misplaced. Thus, individuals may well assume that such AVs possess much more intelligence and operational capacity than is actually the case. Such assumptions may prove critical, if not fatal.

Some of the most evident proximal impacts will be on jobs and associated commuting patterns. The driverless car has the potential to make its human controller as redundant as the horse became for the horseless carriage. Truck and taxi drivers may well have to find new forms of employment, some perhaps supervising these individual vehicles from remote control call centers. However, jobs in the new transportation sector may well

This paper results from the Arthur M. Sackler Colloquium of the National Academy of Sciences, "The Science of Science Communication III" held November 16–17, 2017, at the National Academy of Sciences in Washington, DC. The complete program and audio files of most presentations are available on the NAS Web site at www.nasonline.org/Science_Communication_III.

Author contributions: P.A.H., I.N., and J.S. wrote the paper.

The authors declare no conflict of interest.

This article is a PNAS Direct Submission.

[1]To whom correspondence should be addressed. Email: peter.hancock@ucf.edu.

Published online January 14, 2019.

diminish, as they have in other sectors radically changed by automation and now emerging machine autonomy (8). It is true that some jobs will be created, e.g., in maintaining such fleets of autonomous vehicles, and access to employment for those in economically depressed regions could be improved with AV transport services. Studies (9–11) show that, in general, societal changes resulting from the introduction of these innovations are likely to be extensive. Of course, it is likely that many human-driven vehicles will remain on the roadways for some decades to come. For those who still choose to own their own vehicle, that vehicle need not be parked and taking up space for 22 hours a day. It could be out earning money by giving rides to others. City transit agencies need to consider the arrival of the driverless car now, when plans for future transit projects are in the pipeline. Does a costly subway extension still make sense in light of these emerging transport options? Driverless cars can provide mobility to those who cannot physically drive, such as children, the disabled, or the frail elderly. However, for such populations the problems of ingress into and egress from the vehicle remain, emphasizing that mobility is more than just the car journey alone. Fuller, augmented mobility is a social amenity that can prevent the loneliness, depression, and failing quality of life that attend isolation and immobility. Perhaps such AVs will lure passengers off buses, deleteriously impacting the economics of bus operations in urban areas. These represent only a limited set of the foreseeable changes; more widespread and radical change is promised.

Such enormous changes, when coupled with a still relatively limited public awareness resulting from the somewhat constrained distribution of the details of the precise nature of the technology's development, means that it is important for all involved professionals to disseminate their work and clearly articulate the limits of their own research. Of course, with such complex issues considerable uncertainty remains about the manner in which such innovations will make both short- and long-term impacts (12). There are also further limits on science, and technology communication is itself changing rapidly, especially with the rise of the various forms of social media. Efforts to provide greater opportunities for public scrutiny of science now include avenues for scientists to publish in more popular and accessible outlets (13, 14). It may be that the days of the formal scientific journal are in relative decline. As the progressive forces of speed and utilitarianism affect the processes of research, the need to frame results in terms that the public can comprehend grows accordingly. For example, Google search terms for autonomous vehicles spike when bad news is published, such as following Uber's fatal collision between a test vehicle and a pedestrian in March 2018 (15). The public hears details regarding accidents, but these details are not balanced by more in-depth communications regarding the underlying technical causes or the systems-level advantages of self-driving technology in particular cases and overall. These advantageous aspects include such dimensions as energy savings, overall traffic flow efficiency, mobility for the disabled, and improvements in social cohesion, to name only a few. Indeed, a recent report has observed that the general reaction to autonomous vehicle crashes is likely to be over-emphasized compared with the reaction to collisions in human-operated vehicles and that recent survey data indicate that confidence in AVs is actually slipping (16).

The recent fatal Uber crash came at a critical time for the nascent self-driving vehicle sector. Less-than-perfect cars are being sent onto the roads by companies that have spent billions on research and development and are betting on their success. Uber, Waymo, and others are conducting tests in Arizona, where regulators have taken a hands-off approach to autonomous vehicles; the public there and in other states is being subjected to mass testing without possessing sufficient background contextual knowledge to understand the risks and benefits of such a public experiment. The consent of the public is largely indirect and implicit, since legislative deliberations on these technologies rarely access a public referendum. One of the few ways in which the general public presently encounters such vehicles is through the lens of legal proceedings associated with point failures (17). However, this emphasis may create inappropriate public perceptions regarding the safety dimension of these vehicles (13).

An informed society is important if the public is going to make rational choices regarding cars with no human oversight. Such choices must be informed by scientific understanding of the concerns at issue and must be clearly communicated to a concerned public (18). For example, exactly how safe will such technologies have to be? In 2016 37,461 people died in motor vehicle crashes on the roads of the United States (19). Should AVs be twice as safe, which would mean they kill only 18,000 people per year, or should they be 10 times as safe? Does it matter who specifically is killed, or is this simply a matter of absolute numbers? Who should be regarded as responsible in conditions between AVs and human-controlled vehicles?

The engineers developing driverless cars must consider weighty moral questions. Typically, these are considered as variations on the ethical thought experiment known as the "trolley problem." Should a driverless car swerve and injure one pedestrian if the alternative is to continue straight ahead and injure greater numbers? Should a driverless car protect the occupant above all else or sacrifice the human on board for the greater good in such circumstances? Most specifically, exactly how do we codify these respective ethical and moral principles into a software assembly often created by multiple designers and code developers? We believe that such problems are much more complex than the simple dichotomy expressed in the trolley problem. It is virtually certain that the resolution of such issues must go well beyond the structure of the programming itself to achieve full public acceptance (20). A study that posed such questions to several hundred workers via Amazon's Mechanical Turk service showed that the public remains conflicted about such issues (21). People are in favor of cars that sacrifice the passenger to save other people but would not want to be passengers in a car programmed in that way. These choices are the types of scenario that engineers consider when designing their artificially intelligent machines, but a determinative public discourse and communication surrounding these issues remains lacking. Of course, at this moment most people still lack direct experience of AVs. Risk research suggests that experience and information will lead to people starting to accept AVs, but this propensity does not always hold in all contexts (20, 22). Such acceptance is contingent upon the swirl of public opinion, and such views are founded on how people experience their own interactions with essentially all forms of current technology.

Humans and Machines

As perhaps the primary conduit of the physical expression of human freedom and certainly freedom of movement, transportation in its various forms plays a critical role in virtually all human societies. Mobility is arguably a human right, and when access to such a facility is diminished or denied, the associated quality of life can suffer significantly (23). Especially in the larger and later-developing nations of the world, the availability of ready transport has shaped the fabric, infrastructure, and, to some degree, even the culture of whole countries. Now the very nature of such transportation is changing (24). Since the time of the camel, the donkey, and the horse, humans have occupied the seat of control (25). Humans both have decided on the strategic mission (i.e., the desired destination) and have exercised tactical command (i.e., the momentary control of the animal or vehicle). Of course, this has not been ubiquitously the case. We have always had systems of transport in which a single driver, pilot, or captain exercises control while many others, sometimes numbering in the thousands, are passengers who simply have a passive role rather than any form of active control. There have been rudimentary forms of automation in many transport realms almost from the very beginnings of each technology. Such degrees of automation have increased in sophistication across the decades. Many segments of the transportation system have evolved in this manner with automation becoming an appreciable portion of their technical delivery. Now, however, we are creating and implementing forms of technology that intend to relieve humans of all momentary tactical control and even threaten to supersede all strategic control as well (26). As such, the present expansion is more a disruptive system change than a gradual and controlled evolution. Because disruptions effect broad change, a myriad of issues concerning the future of transport are embedded in many more broadly based questions about the role and impact of automated and autonomous systems in human life more generally (6, 27).

Today's vehicles are already technology-heavy semiautonomous systems built, in turn, by other high-tech systems engineering and production processes. In some circumstances, the coming AVs might well appear to be robots built by robots, although this is not proving ubiquitously to be the case, since many manufacturers now appear to be rediscovering the advantages of human workers. As capital costs of fabrication replace labor costs in the long arc of technical productivity enhancement, we are witnessing not simply a transfer of momentary vehicle control but also a logistical tail effect in which the production and design of those vehicles also shows shifts in decision-making authority. This latter shift of power nominally pits increasingly capable computer-mediated technologies (28, 29) against humans who are largely restricted by the limitations of their inherent capacities to assimilate ever greater tranches of information as technological, robotic progress continues apparently unabated (30). However, the form of progress that constantly features facets of growing machine superiority is not without its own subtleties and caveats. For example, some vehicle manufacturers have found that the vision of fully automated production proves less efficient and sustainable than one in which humans and machines work alongside each other. Like the operations of the innovative vehicles themselves, this could well imply that the human role is not purely vestigial but rather is one that coevolves along with the nature of the technological innovations involved. Such a proposition argues that the development of automation, rather than being a separation of human and technology, might actually represent an on-going symbiosis between humans and the technology they create (Fig. 1).

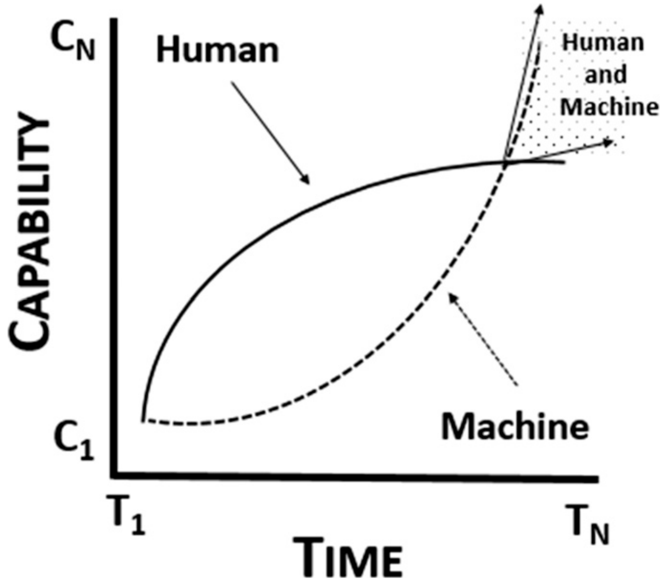

Fig. 1. A contrast between the rates of progress in capability of humans and machines over the recent industrialized epoch. Arguments that technical systems augment human abilities notwithstanding, the generalized curves are revealing. Those arguing for augmentation treat the human–machine system as the unit of analysis. However, it is clear that human and machine still remain fairly differentiated entities, despite efforts to combine them symbiotically. The equivalent start capacities and the exact nature and timing of the crossing point remain highly controversial issues, as does the future path in the combinatorial vector of progress shown in the shaded area of the figure (56, 78). AVs may represent one technology that sees a continuing physical and functional separation between humans and machines.

Previous contrasts of human performance and machine performance have largely been set within the competitive rather than the cooperative perspective. Perhaps the seminal example of such a contrasting comparison derives from the now classic Fitts list (31, 32). In the Fitts report, the authors juxtaposed a series of capabilities in which either the machine or the human excelled and which by implication was performed better by one or the other. This list and those that have subsequently followed it (33) have been characterized as "MABA–MABA" (Machines Are Better At–Men Are Better At) types of analyses (Fig. 2) (34). As indicated, the essence of such comparison is one of contrast, not cooperation. Some have suggested that the very way in which these comparisons are framed is incorrect (35). Technologies can and often do augment inherent human abilities (36); thus rather than an explicit replacement of human knowledge and ability by more efficient machine functions, a partnership between them is sought. Notwithstanding these arguments, two simple curves exhibit the contrasting rates of change in human and machine capability over time (Fig. 1). From this perspective, the relative rates of change mean that machines threaten to supersede increasingly greater swathes of human capacity on an ever-accelerating curve. This eventuality is argued notably by technology optimists who reference increasing computer and robot capability and an associated decrease in the price of their operations (37). However, it must be acknowledged that precisely where we currently stand on these generalized curves has always been a matter of debate (38).

Despite the apparently contrasting trends illustrated in Fig. 1, all is not dissonance between humans and technology, of which AVs are one of the more prominent recent incarnations. For instance, it is clear that only with technology can our limited planetary resources support a world population that soon promises to exceed eight billion persons (39). Generally, the growth of automation does not result in any simple one-for-one replacement of each individual human worker, even within the purview of any particular domain of specific production. Rather the relationship is more complex, with a variegated interplay of augmentation and replacement contingent upon the innovative technology and the fiscal and socio-technical drivers that found such developments (40). In general, automation changes the fundamental nature of the remaining human work. Automation may do this in time increments that cause both benefit and hardship. These effects will be felt for decades and then across generations as economies adapt to the new added capacities (41). Today, these dynamics and influences are reaching most evidently into the multi-trillion dollar social domain of transportation.

One of the proximal problems we now face concerns a quickly approaching and ever greater mix of multiple forms of vehicle control on mutually shared road systems (12, 42). For the foreseeable future many vehicles will be solely and exclusively controlled by human drivers. Indeed, with the enthusiasts' concern for vintage vehicles, such manually controlled machines may continue to be present even in highly automated contexts. However, ever more frequently interspersed with these manually driven vehicles will be systems with some degree of shared human–machine control. Indeed, we are witnessing these systems enter into full-production vehicles at the present. In these vehicles, humans may temporarily drive while the vehicle assumes control on putatively more predictable multilane freeways or on major road arteries, e.g., where traffic contraflow is separated by barriers. The latter driving situation, which overwhelmingly demands precise control over lateral and longitudinal positioning, is especially well suited to current automation capacities (43). However, among this mix of manual and semiautomated control, fully autonomous vehicles will also emerge. This gradual turn to semiautomated and fully automated entities brings particular issues into focus (44). Fully automated systems will be under permanent computer control, never seeking human input and

Fig. 2 drawing and lists
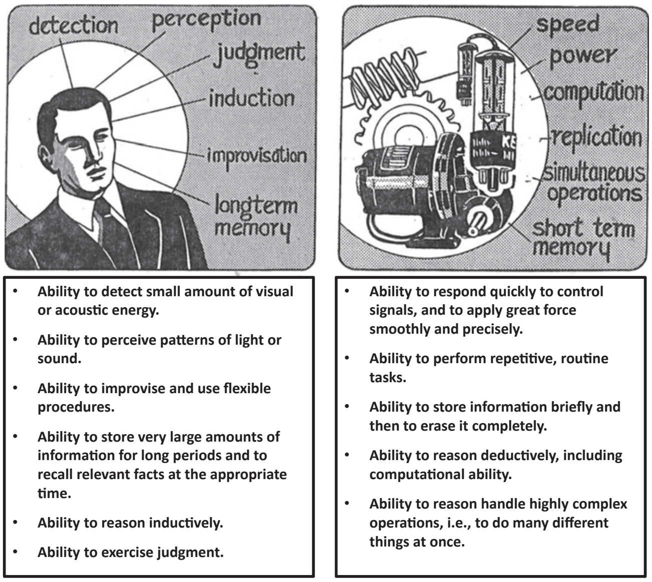

Fig. 2. The comparative lists generated by Fitts and his colleagues (31) in a 1951 report on the future of aviation. Although considered peripheral to the central focus of that report, the lists themselves and the associated graphic have been the subject of much discussion in the more than 60 y of their existence. Most especially, much attention has been paid to whether using direct comparisons provides the most useful strategy for the development of future human–machine collaboration. Such an approach exerts an important impact even today in proposals such as the SAE multiple levels of AV capacity and associated driver interaction. Reprinted from ref. 31, with permission of the Ohio State University Research Foundation.

thus never formally possessing what we now view as an immediate driver. However, all these vehicular technologies will operate together as a community on common roadways, at least for a number of decades. This has been referred to as the "mixed equipage" or "mixed inventory" transportation state. Such mixtures of entities create their own specific problems and conflicts (45). One particularly relevant concern is the psychological concept of attribution error (46).

Attribution is the individual's ability to realize and recognize the motivations of others (47). Human attributions usually work rather successfully. However, when the attributed motivations are incorrect, they lead to cases of attribution error (48). Traditionally, this error occurs when an individual attributes to human causation what is actually the result of the situation or the environment. For example, a person may have thought that someone else's behavior was influenced by a particular motive when, in fact, environmental factors, rather than human factors were the proximate cause of the behavior. We have considerable research concerning when and where people make attribution errors in relation to the motives of others and other motive forces (49). The open question we still face is the precise form and frequency of attribution error when humans act in conjunction with AVs on all segments of the roadway system (50). The roadway is a mutual social resource, so all actors bear shared responsibility for common citizenship (51). In present transportation conditions, such social harmony is preserved and legally mandated by the rules of the road. These rules include formal traffic-control devices and traffic laws and further involve the designed regularities of the roadway by the civil engineers responsible for their construction. Most critically, these rules are also augmented by common assumptions about how other drivers will behave in various driving contexts. While the design features of AVs will obviously evolve during the coming years, the degree to which such features conform to intrinsic human attributions concerning social behavior on the road is unlikely to be uniform (52). For AVs, such concerns fall in the interstices between design innovation and regulation, as indeed they do for virtually all forms of emerging technology. How to reconcile the speed of legislation and the speed of technological innovation has yet to be resolved.

As a general proposition, human beings are effective in distilling the attributions of others. Indeed, the cohesion of society depends upon rational attribution and common ground assumptions (53). Humans driving on the roads together engage in an implicitly choreographed "dance" in which, even though some prove less than completely sensitive, the general collective functions remarkably well. However, the foundational, implicit rules are not always evident. For example, people in different countries and people who come from different cultures drive by different implicit and sometimes explicit rules; one evident example is the side of the road on which they drive. Such cultural, contextual attributions can be highly problematic for strangers to that region (54). Driverless vehicles are the ultimate strangers in our midst. They currently lack the required etiquette to operate as human beings do (55). Unlike the tourist, who at least can depend upon certain common human assumptions, AVs have highly limited access to the implicit rules of the road and access only the explicit rules with which they have been programed. They have even less access to social conventions at the human–human interaction level. Although these latter behaviors might conceivably be programed into AVs, there seem to be few efforts to do so in current production models. At the same time, driverless vehicles may be overly constrained by the formal rules of required behavior that are encoded within their logic. Thus, they may be "bullied" by some aggressive human drivers. In formal human–automation interaction terms, the affordance structures (56–58) of humans and autonomous vehicles are presently incommensurate with one another. The term "affordance," i.e., the relationship between an individual and the action that individual can take in the world, comes from the realm of ecological psychology. Here, we use affordance to mean what actions the driver may take given the immediate driving environment he or she is faced with (59). For example, two-lane roadways afford overtaking in certain conditions but not when an on-coming truck occupies the other lane.

The dissonance between what the human knows of the driving world and what the machine is programed to do will mean that during the approaching transition period conflicts between human drivers and AVs are virtually inevitable (11, 60). While human drivers (and pedestrians) base their affordances overwhelmingly on vision, this is not necessarily true for AVs, which are informed by light detection and ranging (LIDAR), radio detection and ranging (RADAR), and vision as well as other forms of sensors. These various sensors detect other forms and frequencies of emission and so create a perceptual "world" that can be rather different from that which humans perceive. Of course these AV sensor systems must be fused and integrated with each other such that the automated car "perceives" (i.e., assembles its sources of information in pattern-recognition assemblages) the road in a very different way than humans do. This difference is probably not advisable. This divergence of these respective human and machine affordances, and the associated dissonance of attribution, means that human drivers and automated cars are far from achieving the full degree of integration that is currently advertised on many media outlets. Of course, this issue pertains to all developing systems in which humans and automated systems process highly divergent sources of perceptual input information (61).

This issue of attribution failure/error is almost certainly context specific. For example, when passing on multilane highways, there is evidence that human drivers are unable to distinguish other vehicles as being under either human or automated control. Initially, this might seem to indicate successful AV integration: Since people are unable to distinguish between the two, the AV essentially passes at least the surface level of the Turing test in this particular context. However, in such cases indistinguishability can lead to dangerous attribution errors; e.g., humans could wrongly assume that an overtaken car has machine-level perception when, in fact, it is driven by a human who has blind spots (62). These dissonant attributions and their associated errors are liable to be a problematic source of conflict as we proceed down the road to autopia. Thus, such attributions will be highly contingent on the general public's perceptions of AVs, and such perceptions are highly susceptible to the inconsistent coverage provided by contemporary media outlets (63). Of course, attribution errors are not the only issue that beset full human–automation integration.

Technical Feasibility

As is always important regarding a culture-changing technical innovation, historical context can help temper both cynical and overly optimistic reactions to our near future. Research on self-driving cars is not as new as it may seem. Autonomous driving research was well underway in the early 1980s, and by 1996 multiple platforms were demonstrating significant autonomous capabilities, as evidenced by Carnegie Mellon University's "No Hands Across America" project, which demonstrated 96% autonomous driving from Pennsylvania to California along the nation's highways and interstates. The reason these early successes fully two decades ago did not lead to a sea change in driving was not simply technical. As is the case today, the financial model for exactly how autonomous driving would function was uncertain. From sales and lease models to liability and tort, it was unclear how autonomous decision making in machines could be integrated into a legal and jurisprudence system designed entirely for human action and human consequence (64).

A further, more nuanced concern has affected the optimism of technology adopters in this area for the past two decades. The autonomous car can demonstrate statistical success, traveling thousands of miles with reasonable accident rates in the real world. Ironically, the Achilles heel of autonomous machines is the same as their relative advantage: the lack of tactical human decision making. When facing a highly unlikely situation that is nonetheless critical, the machine can face a choice point it has never before addressed: what if a stroller runs into the street after the mother trips and falls; what if two nefarious individuals hijack the car; what if a hacker destroys the braking system; what if glare from a solar reflection blinds the sensors just as a truck passes? Furthermore, we, as human designers and users, may have absolutely no idea how that machine might respond to such novel scenarios. Humans have an empathetic understanding of how other humans behave when facing boundary circumstances. We can imagine ourselves in each scenario, and we can imagine bounds on just how we ourselves might react, but autonomous cars currently express no significant forethought. At present, such vehicles cannot provide hand gestures, they have no bodies with which to jump out of the car and coax the kitten across the street, and there are few prospects for such "individuation" in current and envisaged production vehicles (65). AVs are, at present, sufficiently foreign that the human exercise of empathy toward them fails us: We cannot predict how they will respond to the unpredictable, and therein lies a social science challenge. How do we consider the social ramifications of objects that will pervade society when we cannot even imagine how these interactive, autonomous objects will respond to the boundary conditions that will unquestionably emerge time after time? Rare events are by definition rare, but a one-in-a-million likelihood event will happen millions of times per year if our streets are filled with self-driving machines. This rarity by frequency principle is true in many domains of engineering as well as in human behavior in general; if there are enough propositions, eventually there must be a disposition (66); if we succeed, the unlikely will happen frequently. Engineers have not yet found and may not be able to specify ways to characterize how our autonomous machines respond to the unlikely and the pathological.

A second fundamental concern stems not from any form of direct public antipathy but from another fundamental aspect of statistics and robotics. Robotic devices integrated into the real world depend critically on their sensors and actuators; these are the devices that give them a tangible embeddedness in our world. However, robot sensors and actuators are not human in their resilience to contextual change. Rather tritely, face-recognition software succeeds well, except when it fails. If there are issues of glare, backlighting, sunglasses, blur, or even busy backgrounds, then face recognition often proves unreliable. Fundamental to autonomy is the need to preserve operational success even as context changes and the assimilated information degrades in quality and utility; however, the range of environmental changes that a self-driving car can encounter are extreme. Such real-world challenges are often set in juxtaposition to the laboratory settings in today's research parks and university laboratories where so many autonomy sensors, computational algorithms, and actuators are born. Autonomous cars have indeed driven millions of miles on California and Nevada highways. What, however, does this really tell us about the suitability of driverless algorithms and hardware for national deployment? Road systems vary dramatically across the United States, as well as in other counties around the world. Furthermore, weather conditions in the rest of the United States are significantly worse for the sensors that depend especially on visual understanding of the environment. Heavy rain or snow and poor road markings cause major damage to the vision algorithms in numerous cars today.

Even worse we humans, as designers, engineers, and consumers, can fall short in "engineering empathy" or "technological attribution" in understanding just how and when weather, lighting, topology, and urban clutter threaten autonomy. As long as autonomy is imperfect, the challenge to the human–systems integration remains. Human occupants in driverless cars under this imperative will retain a role similar to that of the ultimate pilot-in-command because their intervention will eventually be required at some juncture as automation fails. It may be that human intervention need not necessarily be from inside the vehicle; instead, as in emerging drone technology, the physical location of the human controller can be remote from the actual vehicle itself (67). Regardless of the specific spatial relations between controller and vehicle, the human operator will require levels of effective situation awareness calibrated to ambient environmental demands at all times because the point at which any such human intervention will be required remains unpredictable as yet. Again, this raises the specter of prolonged vigilance and its well-known decrement and response failure (68, 69). This human-as-backup architecture, which removes the person from momentary control and instead places him or her in a supervisory context, in many ways defeats the very idea of automation in the first place. While this might be seen as a necessary transition phase (6), it will be important to move to full automation quickly, so that the public's expectation of hands-free and responsibility-free personal transportation is fulfilled. Retaining the human in a state of momentary readiness to extract the failing automation from its shortfalls is liable to be not merely unpopular but a major reason for rejecting some forms of shared control in the evolution of AVs.

Of course, it may be possible to engage differing forms of "smart" infrastructure to share these control burdens, but this raises the difficult issue of what parts of the transport system are publicly owned and supported versus the private vehicles that benefit from such social investments. Such discussions are themselves framed within the wider context of an emerging, integrated origin-to-destination–oriented transport system (70). Suffice it to note that many sources of brittleness in the technical system require attention as momentary control passes beyond human hands, if indeed that is the final design goal of these advancing systems.

Autopilot control systems in both private and commercial aircraft provide an informative comparison with ground AVs, but there are some significant differences between the contexts. First, in general, airplanes in the sky face circumstances that present somewhat fewer moment-to-moment dynamic changes and a less diverse range of such challenges than those faced by cars on terra firma. Second, commercial aircraft still have not one but two humans ready to respond with what is anticipated to be high levels of situation awareness. It is clearly now feasible to make a fully autonomous commercial aircraft, an airliner drone.

However, the cost of doing so and demonstrating its absolute safety may well be prohibitive in our current risk-averse climate. As systems become increasingly more complex, our ability to test their responses to all possible states diminishes accordingly, and we quickly reach a point at which it becomes literally impossible to exhaustively test the system as a whole. We then are left with an alternative: to conduct a restricted set of sample-based testing, trying to derive general conclusions from partial evaluation coverage and create resilience by keeping at least one or more humans in the loop (71). The current belief that a human or a team of humans would be capable of innovative, on-the-job problem solving remains persuasive. Humans are not so obviously constrained by the rules written into formal software or by the purely mechanistic limitations of the sensors and effectors of contemporary hardware. Thus it is more than probable that during any extended transition period our semiautomated ground vehicles can and will behave very much like contemporary airborne systems. Whether we will need graded licenses to operate these varying forms of automated has yet to be determined. Precedent argues against such a course.

Of course, this shared control and shared responsibility type of activity is only one of many forms of operation that will be present at any one time. Presently, Uber states that it will field a team of remote human drivers, while Waymo claims to have no need of such support services. Whether shared responsibility is an obligatory phase of development or whether a significant portion of ground transportation can jump to full automation is contingent upon many forces beyond the technical alone (e.g., financial/marketing constraints). The degree to which technological evolution will proceed via gradual change or jump to a new stable state (punctate equilibrium) (72) is still uncertain. Of course, one form of technology may exhibit gradualism while another allied technology vaults forward. That being so, perhaps we should expect a chequered pattern of development of AVs and their on-road operations. It is unlikely that the failure of one sort of AV would prejudice public opinion against all other forms of AV. In comparable terms, the failure of one company's personal digital assistant does not spread immediately to all versions of smart phones or similar, allied technologies.

Examples of boundary conditions can often remind us of the unique value of human innovativeness. Ironically, it is human ingenuity that can give us the examples that require a human common-sense response. A relevant and recent example comes from the computer vision research team that demonstrated the "45 mph stop sign." They created small graffiti modifications of a standard octagonal red stop sign and showed that a remarkably small number of black stickers placed in precise positions on a stop sign accomplish two outcomes simultaneously: to humans, the stickers look like simple graffiti and in no way diminish the ability to discern the stop sign, but to the computer's vision system (as determined by researchers' tests derived from self-driving car software) the very same modified stop sign appears reliably to be a 45-mile-per-hour speed-limit sign. This example is crucial because we are surprised by the computer's mistake and because, as humans, we expect errors to be small deviations from our own norm, not large reversals of interpretation. Collectively we may be surprised if we assume that machine errors follow human patterns of error. They definitively do not. Rather, we must see autonomous machines as true social aliens. They are responsive agents in our universe, but they are nonhuman. They will succeed in ways that are not human, and they also will fail in ways that have nothing to do with how we fail. This prevents us, as nonexperts, from intuitively placing boundaries on the ways in which we expect autonomy to fail gracefully.

Therein lies the grand social challenge of our time vis-a-vis autonomy. AVs may statistically perform more safely than humanity, but are we as humans prepared to live in a world in which failure and degradation appear, from our human perspective, to be random and extreme compared with our comparative expectations of human fallibilities? That is a challenge for systems engineering, for engineering design, and especially for the science of human–machine integration. Progress, especially in this last-mentioned area of inquiry, is not keeping pace with that in the field of computer science or with computational capacity in general. The fundamental problem here is the complexity of understanding human cognition itself. Hence, our desire to innovate responsibly will demand a renewed focus on understanding the societal ramifications of innovation—innovation that is fast but whose integration into society requires us to resolve issues that are not solely engineering challenges but are transdisciplinary social concerns that can be nuanced, complex, and comparatively slow in developing.

Routes to Resolution

In the prior sections, we discussed some of the current and forthcoming challenges that face the integration of driverless vehicles into modern transportation. Having identified such problems, it is incumbent upon us to offer some avenues for research and public engagement that can help us identify and implement more robust deployments. Two vital elements here concern calibrated operator trust and communicated transparency. For the former, design processes should seek to design explicitly for appropriate levels of trust by human occupants in light of the known reliability of the automation (73, 74). This goal is difficult, but achieving it is critical. It is difficult because we are still finding our way in understanding the contextual reliability of differing forms of automation and semiautomation offered by various manufacturers. It is critical, because if there is insufficient human trust in autonomous and semiautonomous systems, there will be both little usage and chronic misuse (75). The company that establishes the most effective calibration of trust, as a result of the highest perceived and actual reliability of their product, will be best placed in this emerging market, that is, assuming the market does not collapse if there is a catastrophic loss of public trust. Transparency is a property of the machine by which it clearly signals both its near- and long-term intent (76). To a degree, vehicles already possess such capacities in various rudimentary forms. For example, displays inside the vehicle already show the driver various states of engine function, gas tank status, and the like. Externally, there are existing but limited techniques for informing other agents of vehicle intent via displays such as turn signals (77). All of these express intent, albeit currently on a rather low level. Such communications, via increasingly information-rich interactions, will greatly enhance vehicle-to-vehicle and vehicle-to-infrastructure conversation capacities.

Trust is already a central theme in the public debate about driverless transportation. Acceptance of any technology requires trust. Quite often new technologies fail to deliver on their promises, either because of the hyperbolic nature of those promises or because their software systems were insufficiently tested and debugged for rich, real-world interaction (78). However, the vehicle is an order of magnitude different from most electronic orthotics. The failure of a telephone or a computer can be annoying and frustrating, and sometimes such failures do even put lives at risk, but on our roadways human lives are always at risk. So, as the risks associated with a momentary failure of interaction increase in proportion to the incipient threat in the environment, the importance of trust grows in lock step. Given that even present-day production vehicles are partial robots, how people come to employ, evaluate, and appropriately trust their robot vehicles will remain a major factor in their greater public acceptance (79). This is centrally a challenge of design, test, and evaluation in conjunction with increasingly widespread public usage.

The trial of trust is further magnified by the fact that these vehicular robots will become more complex as they attempt to function safely in more challenging future contexts, such as dense urban driving. This complexity, in turn, easily outpaces the knowledge level of many of the stakeholders throughout the transportation community. The speed with which new technological innovations increase the gap between stakeholders' intuition regarding vehicles and the reality of what is fielded in our laboratory cities exacerbates an already critical problem. We cannot expect the appropriate form of regulation and policy-setting leadership that is sorely needed until those in a decision-making role possess the appropriate technological fluency and the time to apply it (80). By "fluency" we mean a combination of common semantic grounds for discourse, an awareness of the state of the art, and the critical inquiry skills required to make sense of this new technology. Even then, autonomous systems are still liable to manifest unforeseen positive and negative societal consequences, especially as these transportation advances are linked with other interactive complex systems in society (e.g., the service industry). Fluency itself is a two-way street in which designers, fabricators, manufacturers, and vendors of these emerging systems need to explicate their products in a way that can be understood by legislators and the public alike. Although vital, trust is by no means the only factor involved in such complex decision-making processes; there have been numerous important studies of these constellations of differing influences on decision making (18, 20). We need a new way in which to integrate expertise with decision-making discourse and to nurture trust (16, 81–83). Expertise must itself be trustworthy—not a technical advisor serving at the pleasure of a corporation but rather outreach providing the technology fluency that enables elected officials, civil servants, and the public to engage in meaningful discourse regarding our shared human–technology future (82). In light of the foregoing arguments, we advocate that such an independent scientific body, allied to current transportation regulatory bodies, be created to address the specific concerns of emerging ground-transportation technologies and most especially the issues associated with AVs. Such a group would still face issues in determining how to communicate technical information for general public scrutiny and why such effective information brokerage is so difficult to fully achieve (84).

Conclusions

The advent of AVs on the busy roadways of the world's nations represents a sea change in the way human transportation is conceived and enacted. We are now witnessing an epoch of significant transition in which active control of the vehicle is being taken from the human driver and placed within the charge of the on-board computer systems themselves (85). This transition is reflected in the disparate strategies of current vehicle production companies. Some of the newer manufacturers view the car as a computer that just happens to have wheels, and they have gotten into the business because they are computer companies. Other, more traditional manufacturers persist with driver-centered conceptions. For these traditional names in the motor industry the human, although augmented by numerous assistive systems, remains at the center of the design architecture. Powerful social forces will determine the outcome of these conflicting visions. Of course, there is no reason that the differing forms of vehicle control cannot exist side-by-side on the road and still function effectively. Such a mixture of vestigial artifacts alongside innovative creations often epitomizes much of technology's progress. As greater concern is focused on end-to-end (i.e., starting point to end destination) functionality, the precise incarnation of each part of the journey is liable to be sublimated to the overall purpose. So the car-to-airport, elevator-to-gate, and flight-to-destination routes each requires purposefully directed movement, but now the journey itself is the thing, not the car, elevator, or aircraft per se. In such an integrated world, the AV will take its place alongside these many other modes of transport (e.g., elevators, escalators, bicycles, and ships, among others) that contribute to the ultimate purpose of movement. In this way, we are witnessing the emergence of an integrated transportation infrastructure, and the questions regarding the automation of any one part of the journey represent only a single facet of this much greater concern.

Much ink and much ire has been expended in championing automation as a flawless answer to the continuing carnage of road traffic injuries and deaths. Although it is a laudable goal within itself, it is doubtful whether such a grandiose claim is actually testable or realizable (11, 86). Similar objections can be raised for the claims of greater efficiency, since the pertinent question here is at what level should this efficiency be measured—at the systemic level, or is the claim advanced for each vehicle and person on the roadway? It may well be that all AVs will adhere to the speed limit, but then how does the anxious husband get his pregnant wife to the emergency room? Questions regarding exceptional social conditions go well beyond algorithmic oversight of lateral and longitudinal control. They illustrate one very stark fact: When we change the face of transportation in the manner that is being proposed, we will change the nature of society itself (7). Whether we are prepared for such a radical evolutionary step remains to be seen.

ACKNOWLEDGMENTS. We thank the three anonymous reviewers whose detailed and extensive comments have served to improve significantly the revised work and Dr. Jessica Cruit, Dr. Kristy Snyder, Dr. Frances Martinez, and Dr. Gabriella Hancock for helpful observations on various early versions of this paper.

1. State of California, Department of Motor Vehicles (2018) Permit holders. Available at https://www.dmv.ca.gov/portal/dmv/detail/vr/autonomous/permit. Accessed October 15, 2018.
2. Wakabayashi D (February 26, 2018) California scraps safety driver rules for self-driving cars. *NY Times*. Available at https://www.nytimes.com/2018/02/26/technology/driverless-cars-california-rules.html. Accessed March 9, 2018.
3. Shoup D (2007) Cruising for parking. *Access* 30:16–22.
4. Marshall A (November 22, 2017) To see the future of cities, watch the curb. Yes, the curb. *Wired*. Available at https://www.wired.com/story/city-planning-curbs. Accessed November 26, 2017.
5. Schoettle B, Sivak M (2014) A survey of public opinion about autonomous and self-driving vehicles in the US, the UK, and Australia (University of Michigan Transport Research Institute, Ann Arbor, MI), Report #2014-21.
6. Smith BW (2013) SAE levels of driving automation (Center for Internet and Society, Stanford Law School). Available at cyberlaw.stanford.edu/blog/2013/12/sae-levels-driving-automation. Accessed January 3, 2014.
7. Hancock PA (2018) Some pitfalls in the promises of automated and autonomous vehicles. *Ergonomics,* 1–31.
8. Hancock PA (2014) Automation: How much is too much? *Ergonomics* 57:449–454.
9. Clements LM, Kockelman KM (2017) Economic effects of automated vehicles. *Transp Res Rec* 2606:106–114.
10. Fagnant DJ, Kockleman KM (2015) Preparing a nation for autonomous vehicles: Opportunities, barriers and policy recommendations. *Transp Res Part A Policy Pract* 77:167–181.
11. Securing America's Future Energy (2018) America's workforce and the self-driving future. Available at https://avworkforce.secureenergy.org/wp-content/uploads/2018/06/SAFE_AV_Policy_Brief.pdf. Accessed August 7, 2018.
12. Milakis D, van Arem B, van Wee B (2017) Policy and society related implications of automated driving: A review of literature and directions for future research. *J Intell Transp Syst* 21:324–348.
13. Hancock PA (February 2, 2018) Are autonomous cars really safer than human drivers? *The Conversation*. Available at https://theconversation.com/are-autonomous-cars-really-safer-than-human-drivers-90202. Accessed February 6, 2018.
14. Hancock PA (April 24, 2018) Self-driving cars and humans face inevitable collisions. *The Conversation*. Available at https://theconversation.com/self-driving-cars-and-humans-face-inevitable-collisions-93774. Accessed April 30, 2018.
15. Grabar H (March 19, 2018) Uber crash in Arizona kills woman in first pedestrian death caused by a self-driving car. Available at https://amp.slate.com/technology/2018/03/uber-crash-kills-woman-in-first-pedestrian-death-caused-by-a-self-driving-car.html. Accessed March 26, 2018.
16. American Automobile Association (AAA) (May 22, 2018) American trust in autonomous vehicles slips. Available at https://newsroom.aaa.com/2018/05/aaa-american-trust-autonomous-vehicles-slips/. Accessed May 26, 2018.
17. Gibbs S (January 24, 2018) GM sued by motorcyclist in first lawsuit to involve autonomous vehicle. *The Guardian*. Available at https://www.theguardian.com/technology/2018/jan/24/general-motors-sued-motorcyclist-first-lawsuit-involve-autonomous-vehicle. Accessed January 30, 2018.
18. Fischhoff B (2013) The sciences of science communication. *Proc Natl Acad Sci USA* 110:14033–14039.
19. Insurance Institute for Highway Safety (IIHS) (2017) Available at https://www.iihs.org/iihs/topics/t/general-statistics/fatalityfacts/overview-of-fatality-facts. Accessed January 18, 2018.
20. National Academies of Sciences, Engineering, and Medicine [NASEM] (2017) *Communicating Science Effectively: A Research Agenda* (National Academies Press, Washington, DC).

21. Bonnefon JF, Shariff A, Rahwan I (2016) The social dilemma of autonomous vehicles. *Science* 352:1573–1576.
22. National Research Council [NRC] (1996) *Understanding Risk: Informing Decisions in a Democratic Society* (National Academies Press, Washington, DC).
23. Metz DH (2000) Mobility of older people and their quality of life. *Transp Policy* 7: 149–152.
24. Hancock PA (2017) *Transports of Delight: How Technology Materializes Human Imagination* (Springer, Cham, Switzerland).
25. Herbst J (2006) *The History of Transportation* (Twenty-First Century Books, Minneapolis).
26. Mindell DA (2002) *Between Human and Machine* (Johns Hopkins Univ Press, Baltimore).
27. Hancock PA (2017) Imposing limits on autonomous systems. *Ergonomics* 60:284–291.
28. Moore GE (1965) Cramming more components onto integrated circuits. *Electronics* 38:114–117.
29. Schaller RR (1997) Moore's law: Past, present and future. *IEEE Spectrum* 34:52–59.
30. Ericsson KA, Hoffman RR, Kozbelt A, Williams AM, eds (2018) *The Cambridge Handbook of Expertise and Expert Performance* (Cambridge Univ Press, Cambridge, UK).
31. Fitts PM, ed (1951) *Human Engineering for an Effective Air Navigation and Traffic Control System* (National Research Council, Washington, DC).
32. Licklider JC (1960) Man-computer symbiosis. *IRE Trans Hum Factors Electron* 1:4–11.
33. Sheridan TB (2002) *Humans and Automation: System Design and Research Issues* (John Wiley and Sons, New York).
34. de Winter JCF, Hancock PA (2015) Reflections on the 1951 Fitts list: Do humans believe now that machines surpass them? *Procedia Manuf* 3:5334–5341.
35. Dekker SWA, Woods DD (2002) MABA-MABA or abracadabra? Progress on human–automation co-ordination. *Cogn Technol Work* 4:240–244.
36. Cummings M (2014) Man versus machine or man + machine. *IEEE Intell Syst* 29:62–69.
37. Kurzweil R (2005) *The Singularity Is Near: When Humans Transcend Biology* (Viking, New York).
38. Teilhard de Chardin P (1965) *The Phenomenon of Man* (Harper & Row Publishers, New York).
39. Wackernagel M, et al. (2002) Tracking the ecological overshoot of the human economy. *Proc Natl Acad Sci USA* 99:9266–9271.
40. Hancock PA, Verwey WB (1997) Fatigue, workload and adaptive driver systems. *Accid Anal Prev* 29:495–506.
41. Parasuraman R, Sheridan TB, Wickens CD (2000) A model for types and levels of human interaction with automation. *IEEE Trans Syst Man Cybern A Syst Hum* 30: 286–297.
42. Huang S, Ren W, Chan SC (2000) Design and performance evaluation of mixed manual and automated control traffic. *IEEE Trans Syst Man Cybern A Syst Hum* 30: 661–673.
43. Flemisch F, et al. (2012) Towards a dynamic balance between humans and automation: Authority, ability, responsibility and control in shared and cooperative control situations. *Cogn Technol Work* 14:3–18.
44. Ge JI, Avedisov SS, He CR, Qin WB, Sadeghpour M (2018) Experimental validation of connected automated vehicle design among human-driven vehicles. *Transp Res Part C Emerg Technol* 91:335–352.
45. Desmond P, Hancock PA, Monette J (1998) Fatigue and automation-induced impairments in simulated driving performance. *Transp Res Rec* 1628:8–14.
46. Fiske ST, Taylor SE (1991) *Social Cognition* (McGraw-Hill, New York).
47. Leyens JP, et al. (2001) Psychological essentialism and the differential attribution of uniquely human emotions to ingroups and outgroups. *Eur J Soc Psychol* 31:395–411.
48. Gilovich T, Eibach R (2001) The fundamental attribution error where it really counts. *Psychol Inq* 12:23–26.
49. Follett KJ, Hess TM (2002) Aging, cognitive complexity, and the fundamental attribution error. *J Gerontol B Psychol Sci Soc Sci* 57:312–323.
50. Brooks R (2017) The self-driving car's people problem. *IEEE Spectrum* 34–37:50–51.
51. Button K (2010) *Transport Economics* (Edward Elgar Publishing, Cheltenham, UK).
52. Thomas A (2017) Autonomous vehicles will not prevent half of real-world crashes. *ITS International* (March/April). Available at www.itsinternational.com/categories/detection-monitoring-machine-vision/features/autonomous-vehicles-will-not-prevent-half-of-real-world-crashes/. Accessed May 20, 2018.
53. Jenson J (2010) *Defining and Measuring Social Cohesion* (Commonwealth Secretariat, London).
54. Greenfield PM (1997) You can't take it with you: Why ability assessments don't cross cultures. *Am Psychol* 52:1115–1124.
55. Hayes C, Miller C, eds (2011) *Human-Computer Etiquette* (Auerbach, CRC Press, Boca Raton, FL), pp 351–362.
56. Gibson JJ (1966) *The Senses Considered as Perceptual Systems* (Praeger, New York); reprinted (1983).
57. Gibson JJ (1979) The theory of affordances. *The Ecological Approach to Visual Perception* (Psychology Press, New York), pp 127–143; reprinted (2015).
58. Stoffregen TA (2003) Affordances as properties of the animal-environment system. *Ecol Psychol* 15:115–134.
59. Lewin K (1936) *Principles of Topological Psychology* (McGraw-Hill, New York).
60. Grossman D (May 30, 2018) Tesla hits cop car while allegedly on autopilot. *Popular Mechanics*. Available at https://www.popularmechanics.com/cars/car-technology/a20963081/tesla-hits-cop-car-while-allegedly-on-autopilot/. Accessed June 6, 2018.
61. Hancock PA (2009) *Mind, Machine and Morality* (Ashgate, Chichester, England).
62. Vanderbilt T (October, 2017) Autonomous cars: How safe is safe enough? When it comes to computer-driven cars, reducing traffic fatalities will be far easier than earning drivers' trust. *Car & Driver*. Available at https://www.caranddriver.com/features/autonomous-cars-how-safe-is-safe-enough-feature. Accessed November 3, 2017.
63. Karsten J, West D (May 1, 2018) The state of self-driving car laws across the U.S. (Brookings Inst). Available at https://www.brookings.edu/blog/techtank/2018/05/01/the-state-of-self-driving-car-laws-across-the-u-s/. Accessed May 6, 2018.
64. Coben L (September 7, 2016) Autonomous vehicles: Where morality meets machinery. *The Legal Intelligencer*. Available at https://www.law.com/thelegalintelligencer/almID/1202766891172/Autonomous-Vehicles-Where-Morality-Meets-Machinery/. Accessed September 12, 2016.
65. Hancock PA, Hancock GM, Warm JS (2009) Individuation: The N=1 revolution. *Theor Issues Ergon Sci* 10:481–488.
66. Ayyub BM, McCuen RH (2011) *Probability, Statistics, and Reliability for Engineers and Scientists* (CRC Press, New York).
67. Riley JM, Endsley MR (2005) Situation awareness in HRI with collaborating remotely piloted vehicles. *Proc Hum Factors Ergon Soc* 49:407–411.
68. Hancock PA (2013) In search of vigilance: The problem of iatrogenically created psychological phenomena. *Am Psychol* 68:97–109.
69. Hancock PA, Warm JS (1989) A dynamic model of stress and sustained attention. *Hum Factors* 31:519–537.
70. Fernald JG (1999) Roads to prosperity? Assessing the link between public capital and productivity. *Am Econ Rev* 89:619–638.
71. Hoffman RR, Hancock PA (2017) Measuring resilience. *Hum Factors* 59:564–581.
72. Gould SJ, Eldredge N (1993) Punctuated equilibrium comes of age. *Nature* 366: 223–227.
73. Lee JD, See KA (2004) Trust in automation: Designing for appropriate reliance. *Hum Factors* 46:50–80.
74. Schaefer KE, Chen JY, Szalma JL, Hancock PA (2016) A meta-analysis of factors influencing the development of trust in automation: Implications for understanding autonomy in future systems. *Hum Factors* 58:377–400.
75. Hancock PA, et al. (2011) A meta-analysis of factors affecting trust in human-robot interaction. *Hum Factors* 53:517–527.
76. Hoff KA, Bashir M (2015) Trust in automation: Integrating empirical evidence on factors that influence trust. *Hum Factors* 57:407–434.
77. Norman D (2014) *Turn Signals Are the Facial Expressions of Automobiles* (Diversion Books, New York).
78. Kautonen T (2008) *Trust and New Technologies: Marketing and Management on the Internet and Mobile Media* (Edward Elgar Publishing, Cheltenham, UK).
79. Hancock PA, Billings DR, Oleson KE (2011) Can you trust your robot? *Ergon Des* 19: 24–29.
80. Hancock PA, Hoffman RR (2015) Keeping up with intelligent technology. *IEEE Intell Syst* 30:62–65.
81. Fischhoff B, Lichtenstein S, Slovic P, Derby SL, Keeney RL (1981) *Acceptable Risk* (Cambridge Univ Press, New York).
82. Fischhoff B, Davis AL (2014) Communicating scientific uncertainty. *Proc Natl Acad Sci USA* 111:13664–13671.
83. Gregory R, et al. (2012) *Structured Decision Making: A Practical Guide to Environmental Management Choices* (John Wiley & Sons, New York).
84. Fischhoff B, Scheufele DA (2014) The science of science communication II. *Proc Natl Acad Sci USA* 111:13583–13584.
85. Hancock PA (1997) *Essays on the Future of Human-Machine Systems* (Banta, Eden Prairie, MN).
86. Loimer H, Guarnieri M (1996) Accidents and acts of God: A history of the terms. *Am J Public Health* 86:101–107.