

Alex John London & David Danks, (in press).
Regulating autonomous vehicles: A policy
proposal. In Proceedings of the 2018 AAAI/ACM
Conference on Artificial Intelligence, Ethics,
and Society.

Regulating Autonomous Vehicles: A Policy Proposal

Alex John London^{1,2} & David Danks^{1,3}

¹-Department of Philosophy; ²-Center for Ethics & Policy; ³-Department of Psychology
Carnegie Mellon University
5000 Forbes Ave., Pittsburgh, PA 15213 USA
ajlondon@andrew.cmu.edu ddanks@cmu.edu

Abstract

The widespread deployment and testing of autonomous vehicles in real-world environments raises key questions about how such systems should be regulated. Much of the current debate presupposes that the regulatory system we currently use for regular vehicles is also appropriate for semi- and fully-autonomous ones. In opposition, we first argue that there are serious challenges to regulating autonomous vehicles using current approaches, due to the nature of both autonomous capabilities (and their connections to operational domains), and also the systems' tasks and surrounding uncertainties. Instead, we argue that vehicles with autonomous capabilities are similar in key respects to drugs and other medical interventions. Thus, we propose (on a "first principles" basis) a dynamic regulatory system with staged approvals and monitoring, analogous to the system used by the U.S. Food & Drug Administration. We provide details about the operation of such a potential system, and conclude by characterizing its benefits, costs, and plausibility.

Introduction

Autonomous vehicles (AVs) capable of learning, inferring, planning, and deciding are rapidly moving from a curiosity to an everyday sight in many cities. Much of the debate and discussion about AVs has centered on technological issues and questions: where do they function properly? what are their limits? and so forth. We instead focus here on *human*-centric concerns and responses. In particular, we develop a "first principles" analysis of AV regulation: how (on conceptual grounds) ought we regulate vehicles with autonomous capabilities (independently of local regulatory or legal traditions)?

Current proposals of regulations and guidelines—legal, scientific, and ethical—for AVs (e.g., Hevelke and Nida-Rümelin, 2015; NHTSA, 2016; Smith, 2014, 2016) largely presuppose that existing regulatory systems for non-autonomous vehicles can be suitably modified for AVs. We argue instead that current regulatory systems are, as a type, unsuited to the task of regulating AVs. Thus, we cannot simply expand or generalize the scope and powers of

existing regulatory agencies, but rather must use a different type of approach. In particular, there are important similarities between our knowledge states for (i) vehicles with autonomous capabilities; and (ii) new drugs and other medical interventions. Both domains involve significant interactions, changing or non-stationary environments, (relative) ignorance about many causal factors, and other parallels. We thus propose a novel regulatory system for AVs, modeled on those for regulating medical interventions.

Section 2 is largely negative: it provides a series of arguments and reasons against the use of current regulatory systems for non-autonomous vehicles. Section 3 is, in contrast, largely positive: we provide the key details for the dynamic, flexible regulatory system that we favor, and draw out the analogy with U.S. systems for regulating and approving novel pharmaceutical agents or other medical interventions. Of course, any regulatory system is potentially a double-edged sword, so Section 4 examines both benefits and drawbacks of our proposal, as well as consideration of its feasibility.

Performance Standards and Autonomy

In AVs, decisions and tasks that were once the sole province of human drivers are carried out by the vehicle itself. The AV would thus typically need to be "licensed" as both a vehicle and a driver, but those approvals traditionally have quite different bases. Vehicle regulation focuses on performance standards for their operation in well-defined contexts. Driver regulation involves licensure requirements based on knowledge of rules of the road, context recognition, and other decision making. This latter license tests human abilities under the assumption that relevant knowledge and navigation of a simple test course are sufficient to assess human abilities that drivers will require for arbitrary real world driving tasks. There is no regulatory system that can assess the ability of the vehicle to achieve tasks performed by human drivers within vague or unknown contexts, while conforming to relevant norms.

Standards for Context Identification

A natural response to these observations would be to try to adapt and update the performance standards currently used for vehicular safety and reliability. The U.S. Department of Transportation appeared to endorse this strategy when it suggested “conducting research to develop and validate new performance metrics, establishing minimum or maximum thresholds for those metrics, developing test procedures and test equipment, and then...incorporate[ing] those metrics, procedures, and tests in new FMVSS [standards].” (NHTSA, 2016, p. 68) That is, the U.S. DOT proposes to use similar types of tests and metrics as for, say, evaluations of braking performance. Of course, these AV tests would be more sophisticated, but they would nonetheless be focused on Go/No Go metrics, fixed contexts, and acceptable tolerances on performance.

The problem with this approach, as noted by Danks and London (2017), is that these types of performance standards presuppose that the technology operates within a well-defined *context of operation*. Given such a context, performance standards then typically specify tolerances that must be achieved for specific functions. In many cases (e.g., engine modules), the contexts of operation are largely determined by the operations of other vehicle components.

In other cases, though, a human ensures that the context of operation is appropriate for the deployment of a vehicle system. Traditional cruise control systems, for example, require the driver to continuously ensure that road and traffic conditions are appropriate for maintaining a fixed speed. Dynamic cruise control systems are capable of adjusting vehicle speed without human input, but even they require the driver to ensure that the current context is appropriate for use, and to thereby initiate the system and monitor its performance as conditions change.

One key way in which AVs move beyond mere automation is by assuming the task of identifying relevant contexts of operation, and then selecting (for the estimated or inferred context) the function or behavior that will advance appropriate goals while conforming to relevant norms, constraints, and desiderata. That is, autonomous systems are designed, almost by definition, for situations in which we cannot specify the context *a priori* and where there is not necessarily a human to do that context identification. We thus lack a necessary precondition for the very specification of a performance standard. In short, not only are there no current performance standards for assessing the ability of AVs to individuate and identify relevant contexts of operation, but it is not even clear that such standards are possible (Danks and London, 2017).

Uncertainty and Context Identification

A central challenge to evaluating the safety and reliability of AVs stems from the uncertainty that currently pervades

the context identification and decision making that they are required to perform, in at least the following ways:

- I. Is the set of contexts of operation for a particular AV a closed and well-defined set, or is it open-ended?
- II. For each possible context of operation, are there features that can be used to identify or individuate it? And if so, how reliable are those features?
- III. How sensitive are these features to changes likely to occur within the same environment (e.g., altered driving or pedestrian behavior, traffic density, etc.)?
- IV. Are these key context identification features environment-specific? How do changes in environment affect correlations between contexts and features?
- V. Are the sensors and other data-input devices capable of reliably detecting all relevant context-determining features across all possible environmental conditions?
- VI. Given AV perception of a particular context of operation and desired function, how reliably does its behavior conform to relevant norms?
- VII. If a vehicle misidentifies the context of operation, how quickly and safely can it correct or otherwise mitigate its behavior?

This list is not intended to be exhaustive, nor do we suggest that these questions are insurmountable for AVs (though many have not yet been answered). But these multiple sources of uncertainty illustrate the depth of the challenge of articulating traditional performance standards, as these presuppose that we have positive answers for all of the questions (see also Koopman and Wagner, 2017). Thus, to the extent that AVs will function outside of clear, well-defined contexts, current regulatory systems for vehicles are not capable of providing credible social assurance of AV safety and reliability.

An Alternative Regulatory Approach

At a high level, our scientific, ethical, and epistemic situation with regards to AVs is very similar to our situation for novel drugs and other medical interventions. In both cases, we are dealing with a new technology that provides novel capabilities, but only as part of a complex, non-stationary environment, where we do not know exactly when or why the technology works, or all of the environmental, user, or system factors that make a substantive causal difference. Of course, there are also differences between the situations, but these similarities suggest that the introduction of AVs onto public roadways might be regulated in a manner analogous to the way that novel medical interventions are approved for, and introduced into, clinical practice.

In the biomedical context, *clinical translation* is the process of using insights from basic science to develop a novel therapeutic intervention for a particular medical condition. The process of translating novel medical interventions

from the bench to the bedside has been described and analyzed in detail by Kimmelman and London (2015). Our view is that AV development and deployment from the laboratory and proving grounds to community roads can be usefully understood as “roadway translation,” analogous to clinical translation, and regulated accordingly.

AV “Ensembles” and Application Spaces

The first step in constructing an alternative approach to regulating AVs is to better articulate the “unit” of translation. In particular, we contend that the unit of translation is not any particular piece of software or hardware, or even a software-hardware combination. Rather, these are the most visible elements of a larger “ensemble” of elements that are refined and developed over the course of translation because they are believed to be sufficient to achieve a concrete performance target (Kimmelman & London, 2015). Specifically, an *AV ensemble* includes diverse elements:

- Hardware components (including sensor packages, processing platforms and automotive hardware);
- Software systems (including visual recognition, planning and control algorithms); and
- Models and constraints that specify the space of contexts the vehicle can be expected to identify.

These models and constraints themselves have multiple elements reflecting the complexity of the situation:

- A set of environments in which those contexts must be identified;
- Degree of environmental and contextual variation with which the system must be expected to cope; and
- Behaviors that the AV should be expected to perform, along with context → behavior mappings.

The task of AV development is then to assemble a combination of technologies, along with models and constraints that specify the conditions under which developers believe the AV will be a safe and reliable technology. The actual real-world safety and reliability of an AV ensemble thus depends partly on: the likelihood that the vehicle will only be deployed within relevant limits; whether the model and constraints adequately capture real-world operational demands; and the ability of the various systems to safely and reliably perform those tasks under real-world conditions.

To better understand the non-hardware, non-software elements of an AV ensemble, consider first a delivery truck that is developed only to convey packages from a central distribution point (such as an airport or freight terminal) to a limited number of ancillary distribution points via controlled roadways. In this case, the *AV ensemble* includes limits on contexts of operation, such as the limited number of possible routes for the vehicle, and constraints that are assumed for those routes, such as the nature and kind of obstacles and traffic it might encounter.

Suppose instead that the exact same delivery truck (in terms of hardware and software) is developed with the goal of taking packages from a central distribution point to an *arbitrary* address within a single city. The hardware and software package are now part of a different AV ensemble, precisely because the system is expected to deal with a wider set of contexts with increased variability both within and across contexts. Moreover, the models and constraints used to specify the set of environments in which the AV is expected to operate safely, and the universe of tasks it is capable of performing, must be correspondingly richer.

This characterization of an AV ensemble makes explicit what diligent developers already know: successful design requires matching system capabilities with the task and environmental demands for intended use-cases, including plans for unexpected contingencies, and explicit specification of what counts as success in various situations. As most developers recognize, sometimes we alter technological aspects of the vehicle (software or hardware), but sometimes we alter the AV ensemble by expanding or adjusting the deployment space. This latter approach is not available with traditional vehicles, as they have a well-defined deployment space. Any successful regulatory or oversight system for AVs must thus take a broader view of the unit of translation: the safety and reliability of AV ensembles needs to be evaluated whenever changes are made to *any* feature of the AV ensemble, not just obvious changes in hardware (e.g., new sensor array) or software (e.g., new code). We thus require a staged process of testing and rollout into different deployment spaces, coupled with ongoing review.

Early-phase Testing

Regulation of a particular AV ensemble should begin with early-phase testing. This stage is exploratory in nature and includes simulation studies in virtual environments, as well as detailed testing of the full (though presumably constrained) AV ensemble in carefully controlled physical environments. Its primary goal is to refine the AV ensemble to locate the combination of elements whose safety and reliability will be evaluated in later-phase testing, and to identify potential failure points. A central part of this refinement process is exploration of the application and deployment space. Because operational contexts cannot be fully specified in advance, early-phase testing should aim to survey plausible contexts of operation so that we have a positive characterization of windows within which the AV can be used and outside of which failure is more likely. Empirical testing is critical in this step; we should not presuppose that developers can specify contexts *a priori*. Moreover, developers may well need to update their understandings of the space of contexts of AV application in the course of exploring the possibilities.

For example, a delivery truck that is expected to travel to only a few destinations will still have to respond to changes in the context of operation such as road closures or other detours. The safe operation of such vehicles may require that either they be capable of navigating arbitrary routes to the same destination, or their acceptable contexts of use may need to be correspondingly restricted.

As the exploration of the scope of application uncovers additional relevant contexts of operation, developers can assess the ability of AVs to identify those contexts and respond with appropriate functions. The use of controlled contexts enables developers to alter conditions systematically to challenge AV systems and identify points of failure, or windows outside of which AV performance begins to degrade. Identification of failure points and windows of reliable functioning is critical to refining all components of AV ensembles. It is also essential to ensuring that AVs are capable of functioning safely and reliably under the less strenuous conditions of routine deployment.

A goal of early-phase testing should be the development of AV system models that codify conditions under which the system has been tested. Those models can then provide predictions about which environmental or contextual variations are likely to degrade or breakdown system performance, and when it will function reliably. Through iterated simulation and controlled testing, these models and constraints can be tested and refined.

An additional deployment goal for real-world settings is estimation of the extent to which vehicles encounter situations that are expected or unexpected relative to the models and constraints used to predict their performance. Unexpected events will likely be correlated with system failures. As changes are made to avert such failures, we will need to continue to recreate a range of situations (including failure cases) so that the system performance in permutations of such situations can be assessed.

Before an AV system can be introduced into real-world contexts, it should be possible to provide regulators with compendious information about each of the elements in the AV ensemble. Regulators obviously must know the hardware and software, but they also must be informed about the models and constraints that guided system development, and now ground performance expectations.

Transitional Testing

Once developers believe they have refined an AV ensemble to the point that it is capable of identifying and responding to anticipated contexts of operation, then testing should move into carefully monitored deployment in real-world settings. Initially, the goal of this stage of testing is to gain real-world experience with an AV ensemble so that we can assess its performance at identifying and responding to contexts of operation that were anticipated in the

models and constraints that both guided the AV's development and also shape expectations about reliability under broad, real-world conditions.

At the same time, real-world settings are likely to be more complex, more variable, and contain other unanticipated features, so testing AV ensembles in such settings will enable developers to uncover additional, unanticipated contexts of operation and also refine the AV elements. Throughout, AV performance must be carefully monitored and controlled since we initially lack real-world experience. It should also begin with "ideal" conditions and then move into more difficult scenarios, progressively challenging AVs with more complexity, noise, and novelty.

This transitional testing is importantly *not* simply "deploy the AV on public roads and record what happens," just as drugs do not move immediately from early-phase testing to widespread prescription and monitoring. Rather, transitional testing involves careful design of test activities to specifically examine performance in particular contexts or conditions. For our regulatory structure, claims such as "this AV has operated accident-free for two million miles" are relatively uninformative, as they provide no details about the contexts or deployment spaces in which those miles have occurred (see also Cummings, 2017). Of course, such measures are not *completely* uninformative, as the AV likely did encounter several different contexts across that many miles. But regulatory purposes require performance evidence across a range of contexts and situations, and that requires careful thought and design to ensure that the AV ends up in those positions.

When AV ensembles are designed to operate in highly controlled spaces of application, it may be easier to discover the full range of contexts of operation that the AV will subsequently encounter during normal operation. As expected deployment spaces become more open-ended, it becomes correspondingly more difficult to anticipate novel contexts of operation. As a result, greater weight will have to be placed on the probability that the models and constraints that guided system development, and that ground predictions of its reliability, are adequate to capture the demands of real-world environments. The key is to provide reasonable public assurance that AVs identify and reliably function in contexts of operation that can plausibly be anticipated to arise.

This period of transitional testing ends when developers believe, and the oversight agency agrees, that their AV ensemble has been sufficiently refined that it is capable of operating independently in the designated space of application. This regulatory decision-making will obviously be quite tricky in certain cases, but does not pose a new type of challenge for a regulatory entity. Rather, it can be addressed using processes similar to those already used by various agencies to set criteria and thresholds (e.g., deliberation, public comment, and so forth).

Confirmatory Testing

The goal of confirmatory testing is to gather data on AV ensemble performance under closely monitored, real-world conditions with two goals. First, we aim to validate the hypothesis that the AV ensemble is capable of detecting and responding to the range of contexts of operation for the specified space of application. Second, we aim to quantify the reliability of the AV ensemble's performance in uncontrolled conditions over an adequate period of time.

Initially, confirmatory testing should be conducted under the supervision of trained engineers, where interventions into the AV operation are logged and examined to determine the source of problematic behavior. Adjustments can then be made to address the issue. Further transitional or confirmatory testing will often be required to validate the efficacy of those AV ensemble changes.

In many ways, confirmatory testing is analogous to the prescription-and-monitoring stage of drug development. In the latter case, drugs are prescribed only by specially trained individuals (doctors) who are educated in the risks, benefits, and appropriate contexts for those interventions. Moreover, follow-up monitoring of a drug typically occurs for many years after it is first approved for prescription (for particular conditions), precisely to track whether unexpected symptoms or side effects occur once the drugs are used in a wider variety of contexts.

Similarly, this confirmatory testing for AV ensembles will provide engineers and the oversight/regulatory entity with information about real-world performance in less-controlled settings. Just as this type of information must be provided by drug manufacturers even after initial approval, it is critical that the regulatory entity for AVs have access to this type of performance data for the whole ensemble (not just the technology). Traditional vehicle regulation does not need to know how the vehicle passes the test, only that it successfully meets the performance requirements. In contrast, AV regulation, whether in the form of certification or licensing, must be based on more fine-grained information about the performance of the whole ensemble (software-hardware technology, plus models and constraints that ground projections of reliability in different contexts of operation), analogous to the fine-grained information required by drug approval agencies.

Regulatory Entity

Throughout this section, we have repeatedly referred to some unspecified "oversight or regulatory agency." The alternative regulatory paradigm we describe here would require an entity analogous to the U.S. Food and Drug Administration to control access to consumer markets and to ensure safety and reliability. This regulatory entity (RE) would need to be empowered to (a) restrict the deployment and sale of AVs until evidentiary benchmarks for safety

and reliability can be demonstrated; (b) monitor post-approval performance information; and (c) require modifications to approved AVs.

To discharge its mandate, the RE would require the ability to compel disclosure of all data relating to an AV ensemble from all stages of design, development, and testing. Because these data will undoubtedly include information about proprietary systems, the RE would need mechanisms to assure confidentiality of such information.

Working with relevant stakeholders, including transparent processes of public engagement, the RE would also be tasked to articulate context-sensitive benchmarks for safety and reliability that will provide initial thresholds for market approval. For example, these may first take the form of population-level achievements, such as a certain number of driving-hours in various contexts without human intervention or tolerable accident rates. Of course, approval would be restricted to those contexts in which the AV ensemble achieved those benchmarks.

Over time, the RE will presumably build significant expertise about relevant contextual dimensions, challenging situations, and so forth. This information could readily be passed back into the testing processes for AV ensembles. Ideally, this process could lead to the creation of regulator-designed simulation environments and real-world tests that can be shared. The long-run upshot would be the development of significant expertise in a "trusted partner" that is not itself in the industry.

Pros, Cons, and Realism

This regulatory system is clearly more complex than we typically find for vehicles, which carries significant benefits, drawbacks, and complications. The most obvious benefit is, of course, increased confidence that public safety would be improved through the use of AVs. A number of different positive benefits of AVs have been proposed (e.g., Anderson *et al.*, 2016; Fagnant and Kockelman, 2014), but we do not yet have the relevant real-world data to judge those claims.

A related benefit would be increased trust in the AV ensembles themselves, whether by users or the public that shares the roadways with them. In general, the introduction of autonomous capabilities into a system makes it harder to develop trust in that system (Roff and Danks, in press), and this problem is exacerbated in the case of AVs. Lack of trust is a key driver of non-adoption of a technology (e.g., Gefen *et al.*, 2003; Pavlou, 2003), and the benefits of AVs obviously cannot be realized if they either are not used, or are highly legally constrained because of lack of public trust. The regulatory system that we propose here would ensure the disclosure—at least, to relevant representatives of the public interest—of exactly the types of information

(roughly, why an AV performs as it does) that are required for the development of the requisite types of trust (Roff & Danks, in press).

The previous benefit also indirectly supports AV developers, as it increases the likelihood that their systems would (if successful) be widely adopted. It rewards AV developers who create high-quality, reliable products by excluding (from the market) developers who sell similar “looking” systems that are unable to perform adequately. Lower cost products would potentially capture greater market share because of their attractive price, but also endanger the market as a whole since higher rates of accidents or fatalities may breed distrust of the entire AV market, rather than being limited to a few bad apples.

Finally, the information acquired by the RE could, due to its status as a trusted partner, significantly advance the state of the art in AV development. Exchange of information between AV developers would presumably benefit all, but such exchanges are highly improbable for competitive reasons. An independent, trusted RE could facilitate information dissemination and transmission among the various AV developers.

Of course, our proposed RE achieves these benefits at the cost of being more invasive than current vehicular regulatory agencies (e.g., the U.S. DOT). Increased testing and disclosure requirements would likely slow some development. At the same time, we frequently decide as a society that such tradeoffs are worthwhile. For example, standards for the marketing and sale of drugs are high because of the high costs—squandered financial resources and human health—of allowing firms to sell ineffective or harmful interventions. Moreover, because of the complexity of medical technology, we cannot simply inform consumers and trust that they will make an optimal decision. AVs and drugs are not like cell phones, where unreliability can be punished in the marketplace, since the consequences of AV failure are likely to be more significant (but less predictable by consumers) in terms of injury and death. As a result, the marginal costs in speed of development are likely more-than-outweighed by the benefits of this proposed regulatory system.

A second concern is an ethical one: this type of dynamic, staged regulatory system inevitably involves a certain amount of discovery through testing on the public. Moreover, unlike clinical drug trials where participants can provide informed consent, the use of AVs on public roadways will potentially affect many individuals who never gave such consent (Cummings, 2017). While early-phase and transitional testing can provide clarity about the performance contexts and characteristics of an AV ensemble, we will never have complete information. After all, the whole point of the confirmatory testing phase is exactly that we expect real-world deployment to involve contexts and conditions for which we do not have relevant performance

data. But as a result, this phase necessarily includes an element of continued testing, now on a public (e.g., other drivers) who did not explicitly consent to participation in this particular “experiment.”

We agree that this is a legitimate ethical concern, but we also suggest that our proposed regulatory system somewhat mitigates it, and that it is certainly superior to current systems. Our proposed RE would have explicit regulatory *and monitoring* authority, and so would be able to intervene on the public’s behalf if the real-world “experiment” goes awry in some way. Our previous arguments imply that some degree of real-world testing and learning must occur; we cannot do everything in the laboratory. The ethically best response is to embrace this fact and design our regulatory systems to ensure that, for example, the “experiments” can be adjusted or halted in a timely fashion if necessary, as our proposed RE would be able to intervene.

We have largely sidestepped issues of political feasibility here. We expect that there would likely be significant barriers to establishing such an RE, as expected benefits are diffused across a large number of people, while the expected costs are localized on a small group (i.e., the AV ensemble developers). Moreover, there would be natural worries about regulatory capture in such an RE. Nonetheless, we believe that this “first principles” analysis of the regulatory system that we ought to prefer provides a relatively concrete target for future debates and discussions.

Conclusions

AVs are rapidly spreading, and debates about how best to regulate them currently lag far beyond the technology. Current regulatory systems and agencies for vehicles depend on clear benchmarks for well-defined contexts. Autonomous systems are valuable, however, precisely when the contexts are vague or underspecified, or where it is not necessarily clear what counts as “success.” We thus must shift our understanding of how to regulate these systems. The introduction of autonomy is not similar to the introduction of a novel braking system, but instead is a completely different type of feature.

Rather than focusing solely on the technology, we must broaden our field to the full AV ensemble by including the assumptions and preconditions that are known to the developers, but are typically not publicly discussed or disclosed. Rather than using Go/No Go criteria, we need a dynamic, staged system that gradually increases the approved contexts of application through directed experimentation by informed users. Rather than blanket approval, we need continued monitoring and refinement to ensure public safety and trust. In short, we propose that AV ensemble regulation should resemble the regulation of drugs and other medical interventions.

References

- Anderson, J. M., Nidhi, K., Stanley, K. D., Sorensen, P., Samaras, C., and Oluwatobi, A. Oluwatola. 2016. *Autonomous vehicle technology: A guide for policymakers*. Santa Monica, Calif: RAND Corporation Press.
- Cummings, M. 2017. The Brave New World of Driverless Cars. *TR News* 308: 34-37.
- Danks, D., and London, A. J. 2017. Regulating Autonomous Systems: Beyond Standards. *Intelligent Systems* 32(1): 88-91.
- Fagnant, D. J., and Kockelman, K. M. 2014. The Travel and Environmental Implications of Shared Autonomous Vehicles, Using Agent-based Model Scenarios. *Transportation Research Part C: Emerging Technologies* 40: 1-13.
- Gefen, D., Karahanna, E., and Straub, D. W. 2003. Trust and TAM in Online Shopping: An Integrated Model. *MIS Quarterly* 27(1): 51-90.
- Hevelke, A., and Nida-Rümelin, J. 2015. Responsibility for Crashes of Autonomous Vehicles: An Ethical Analysis. *Science and Engineering Ethics* 21(3): 619-630.
- Kimmelman, J., and London, A. J. 2015. The Structure of Clinical Translation: Efficiency, Information, and Ethics. *Hastings Center Report* 45(2): 27-39.
- Koopman, P., and Wagner, M. 2017. Autonomous Vehicle Safety: An Interdisciplinary Challenge. *IEEE Intelligent Transportation Systems* 9(1): 90-96.
- National Highway Transportation Safety Administration (NHTSA). 2016. *Federal Automated Vehicles Policy: Accelerating the Next Revolution in Roadway Safety*. <https://www.transportation.gov/AV/federal-automated-vehicles-policy-september-2016>
- Pavlou, P. A. 2003. Consumer Acceptance of Electronic Commerce: Integrating Trust and Risk with the Technology Acceptance Model. *International Journal of Electronic Commerce* 7(3): 101-134.
- Roff, H. M., and Danks, D. In press. "Trust but Verify": The Difficulty of Trusting Autonomous Weapons Systems. *Journal of Military Ethics*.
- Smith, B. W. 2014. Automated Vehicles Are Probably Legal in the United States. *Texas A&M Law Review* 1: 411.
- Smith B. W. 2016. Automated Driving Policy. In G. Meyer and S. Beiker (Eds.), *Road Vehicle Automation 3. Lecture Notes in Mobility* (pp. 53-58). Springer, Cham