# Technical Problems in Autonomous Vehicle Policy Research

**Anonymous Author(s)**
Affiliation
Address
email

## Abstract

The deployment of Autonomous Vehicles (AVs) will have a transformative impact on society. These impacts could extend far beyond the safety and efficiency of individual vehicles. By changing how road users interact with the transit system throughout society, AVs could generate large economic side effects and shifts in traffic throughput, leading to new policy questions bearing on multi-modal routing, fair route assignments, and optimized geofencing. Identifying these externalities and designing AVs to address them could be beneficial to all commuters, beyond individual passengers. In this work we present several effects studied by social scientists conducive to technical work, and propose a new subfield to coordinate the identification and modeling of AVs' social impacts between research communities.

## 1 Introduction

Autonomous vehicles (AVs) will have a transformative impact on society and transportation infrastructure. Beyond personal safety and economic activity, AVs may also generate unprecedented externalities, including undesirable discrepancies in physical mobility, economic mobility, and pedestrian detection [1]. While these externalities have been studied by social scientists, we are unaware of technical research which accounts for them from the designer standpoint. One barrier to conducting such work is the lack of concrete technical problems to be addressed. Our goal is to provide these concrete problems. By designing systems which mitigate negative externalities before they can occur, we can better ensure that the macroscopic effects of AVs will be beneficial to everyone.

To avoid each unintended externality of AVs, we must achieve three goals:

- Identify the externality
- Model and control for the externality
- Enact legislative policy to ensure the proper control

For instance, to identify how self driving cars could cause differential impacts on physical mobility, we could model systems which provide more or less equal access to transportation, and then enact policies to ensure that AV systems account for this. We could similarly control for how the behavior of AVs differs depending on the safety concerns and special considerations of communities they are routing through, how the commuting times should be balanced between different communities, and how they will spur or suppress investment into road maintenance. Of course, there is no one community of researchers well suited to tackle all of these goals. Social scientists and behavioral economists are studying relevant externalities, technical AV researchers are well prepared to make AVs which control for these effects, and policy makers will be able to enact these solutions. We are calling for an ongoing collaboration and coordination between these groups to achieve these means.

Submitted to 34th Conference on Neural Information Processing Systems (NeurIPS 2020). Do not distribute.

The major bottleneck for work in technical AV policy is the lack of clear concrete problems, and the goal of this document is to show what technical AV researchers can do to help: model—and control for—these externalities. Technical work in this direction is both tractable and well situated to have a large impact towards making the overall outcome of AVs beneficial. In Section 2 we give a broad overview of the problem space. In Section 3, we identify how some known externalities of AVs can be modeled to produce well-specified technical problems.

# 2 The Space of Externalities

Most work ensuring beneficial impacts of AVs has focused on the safety of the people in and around the vehicle[2, 3, 4], which has the potential to save tens of thousands of lives annually. However, AVs will also reshape transportation, and in turn, have ripple effects to other layers of social interaction. There has been limited technical work in modeling these large scale effects. Work on the effects of local behavior of self-driving cars on large scale traffic patterns has revealed that default approaches to routing will cause large-scale congestion [5]. Once identified, technical solutions were designed to mitigate these effects, which allow for local routing solutions with good aggregate behavior [6]. To mitigate the externality of large scale traffic congestion, the first step was to identify it as an externality, model the effects, and construct some control for those effects. In a world where this effect was not identified, was not modeled successfully, or was not controlled for, this effect would have manifested in systems causing measurable and preventable harm.

While previous technical work on AVs has focused on effects within the transportation system, which we refer to as the *technical effects* described by the left column of Figure 1a, we focus on the effects outside the transportation system, which we refer to broadly as *social effects*. These are effects that changes to transportation have on all of the parts of our society that use it. In local interactions (bottom right quadrant of Figure 1a) there will be dramatic effects on how individuals experience roadways – how pedestrians signal to cross the street, how drivers signal to change lanes, who is blamed when a crash occurs, and the general norms of how people use the roads. Considerations about how assertive or passive cars are today could have long-run effects on the sorts of driving strategies that are successful for the remaining human drivers. There is less work on mitigating effects in this upper right quadrant. AVs' ability to coordinate on large scales will likely have significant effects on road conditions and will actively intervene on physical mobility, safety, and comfort of different communities. These effects are wide ranging but have been a focus of study in the social sciences. As a result, there are many models that have been made of these effects and working to mitigate them is a tractable research problem with the tools available today.

We should expect effects like these to arise from AV interventions. Figure 1b describes one way to get a sense of the scale of these effects, representing all the subsystems involved in transit and imagining how decision-making could change if these were automated. Since transportation is central to social order, interventions would have ripple effects on distinct aspects of society. Thus, introducing AVs is not as simple as making all of the trips that we currently take more efficient and safe. It will affect traffic patterns, which will in turn change the relative ease of getting to different locations and how comfortable and safe the traffic is around us when we get there, which will affect the behavior of individuals and businesses, which in turn will affect traffic patterns. This has the potential to reshape many facets of society, and though we don't fully understand the effects, social scientists have been hard at work modeling what the effects may be. Though many of these effects are beneficial – easier and cheaper access for many to reliable and safe transportation – others are not as clearly good.

However, these effects are not inevitable. As the designers of these systems, we have the ability to anticipate and address these externalities. Moreover, social scientists, by modeling the effects of our actions, have been making the tools necessary to get started.

# 3 Modeling Externalities for Control

Here we give an overview of externalities that have been identified, as well as models which could serve as the basis for technical systems accounting for their effects. Solutions to these problems could be at the level of local routing, at the level of managing traffic patterns, at the level of managing infrastructure, or some combination of these. Though this list is not exhaustive, we hope it serves as a starting point for designing systems which account for these identified externalities.

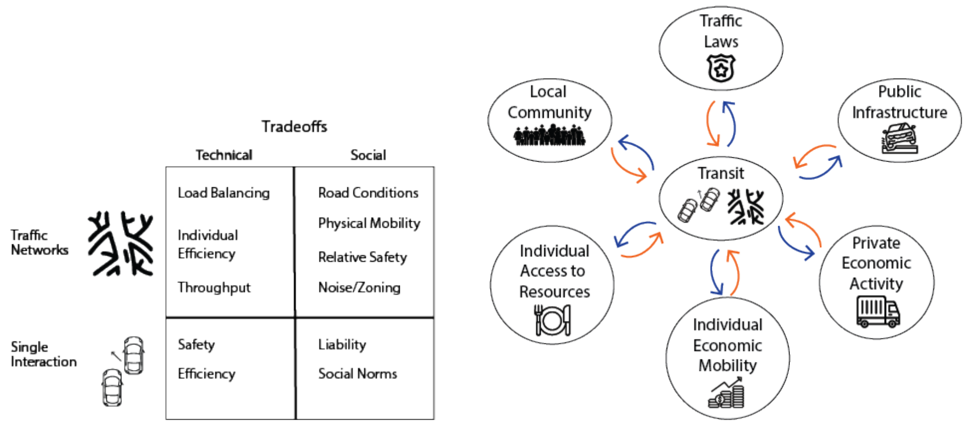


(a) Considerations About the Transit System

(b) Systems with Which the Transit System Interacts

Figure 1: Technical and Social Dimensions of AV Development

### 3.1 Community-Specific Routing and Control

AVs will need to operate in ways that community stakeholders and neighborhood representatives affirm as legitimate and trustworthy, rather than merely safe. This means that the AVs will need to operate differently depending on the human factors considerations of the communities they route through, including emergency preparedness [7], navigation wayfinding [8], and other local concerns.

For example, residential neighborhoods in the United States often accommodate special needs groups through distinct signage: warning signs about pet dogs and cats, "children at play", and protection for the disabled (e.g. audible walk signs). Beyond vehicle features such as wheelchair access, AVs will need to incorporate routing adjustments so that time spent in these zoned areas is minimized. Meanwhile, some communities require unique forms of road mobility, such as retirement facilities and golf courses that have their own specialized modes of transport. Some communities have adopted special guidelines for golf carts interacting with normal traffic vehicles [9]. Each of these considerations, and other details of local customs which we have yet to consider, need to be incorporated into the local control procedure so that they can be customized by the communities to be contextually appropriate.

### 3.2 Fair Congestion Management

Behavioral economists have studied the potential for AVs to affect driver behavior through dynamic congestion pricing, modeling the effect with agent-based models [10]. This method leads to a well-specified set of metrics that can be used as targets for optimization when designing AVs. Concretely, parameters for departure time choice and route choice help capture induced demand and congestion effects on particular stretches of highway. Drivers are modeled as making rational choices under some cost function, while total social welfare could be measured both before and after dynamic tolling. Given this model, it is possible to calculate optimal dynamic tolls and study how quickly the model converges to optimal conditions.

This model provides a way to determine how routing considerations will affect user choice, and how that will affect public revenues and travel times. It is also straightforward to extend this model to understand how these effects would manifest in different communities. This leads to a wide assortment of natural technical problems: designing routing procedures to maximize welfare, minimize congestion, ensure equitable access to mobility across communities, or balance the performance of shared and individually owned AVs. These have the potential to trade-off against other considerations, such as zoning rules, local efficiency, and safety. Considered as a whole, this points to a vast unexplored space of well-defined technical problems whose solutions would help ensure that the benefits of AVs are distributed fairly and effectively.

### 3.3 Differential Community Impacts of Capabilities

While AVs will affect infrastructure and traffic, infrastructure conditions will also limit AV deployment, and there must be control over resulting feedback loops. The extent to which AVs can be deployed will be affected by local road quality, weather patterns [11], and the economic status of residents. If this persists, then the large economic advantages of AVs could be localized to communities with ideal conditions, unless infrastructure or technology is adapted to mitigate these effects.

Moreover, as AVs become widely deployed, their effects and impact on public infrastructure (roads, bridges, highways, electrical grids) may be felt unequally. For example, while AVs must avoid potholes successfully and consistently, there is a tension between modeling potholes in terms of perception (identify them as they appear) or route planning (avoid roads that are more likely to have them). As feature detection of road damage remains a stumbling block, the latter seems more likely for the foreseeable future–and this is likely to generate effects on congestion, highway flow, and other macro-traffic dynamics. AVs can compensate for this by measuring and minimizing loss in fuel efficiency or average time to destination from avoiding potholes, aiming to preserve the road without disrupting traffic. This be quantified through reference to existing models for highway maintenance [12], priority damage assessment [13], and smart pavement evaluations [14]. This work could aid constraint satisfaction by including factors that corroborate existing public standards for road maintenance, rather than modeling vehicle motion in isolation.

Work to mitigate these effects could focus special attention on groups which are disproportionately affected. For instance, human factors research on how and why humans fail to predict vehicle intent [15] could help route optimization in distinct weather environments, as AVs could adjust driving speeds in response to distinct weather features in order to match the expectations of surrounding pedestrians and thus still ensure safe deployment. Even if the effects are difficult to mitigate, they could at least be measured by building tools which assess access to transit between different groups.

### 3.4 Differential Economic Mobility

If developers allow route planning to be completely determined by local constraints, this could result in transit dead zones (i.e. areas that AVs systematically avoid), greatly reducing access to transportation as a whole. The resulting declines in physical mobility would inhibit social mobility by limiting access to jobs and other opportunities. AVs will need to minimize the likelihood of generating transit dead zones at the same time they establish regular connections between urban cores and exurban areas.

Physical mobility has measurable economic effects on individuals, as access to the city center is a central and longstanding concern of urban planning and transportation infrastructure [16]. Macroeconomists have recently expressed interest in AVs' potential for economic disruption via geofencing, comparing their impact scale to the internet and the original creation of the Interstate Highway System. Prominent models have focused on improved access to labor markets and low-income mobility, among other factors [17]. These models serve as a possible metric to quantify the effects of transportation access on an individual basis.

Access to physical mobility itself can be characterized by the latency and cost of ride sharing platforms, which are affected by the overall network congestion, which in turn are affected by the routing behavior of AVs. Thus, it should be possible to determine how changes in the routing procedures of AVs affect the physical mobility of certain communities, and in turn determine downstream effects on economic mobility. This leads naturally to questions about how AV routing could preserve fair access to transportation and economic mobility, or at least finding ways to accurately quantify these effects so they can be monitored as AVs are deployed at scale.

## 4 Conclusion

We have collected several externalities that are ready to be modeled, measured, and prospectively mitigated by AV developers. While a full exposition of all externalities to be generated by AVs is beyond our scope, we envision this document as a first step in an emerging research subfield at the intersection of technical AV design, human factors research, and public policy. We hope this serves as groundwork for potential collaborations both between researchers in human-robot interaction, ML fairness, and AI Safety, as well as social scientists and transportation planners in the long term.

## References


[1] K. Wilson, T. Snyder, A. Schmitt, A. Short, and T. Dutzik, "Study: Avs may not detect darker-skinned pedestrians as often as lighter ones," Jun 2020. [Online]. Available: https://usa.streetsblog.org/2020/06/17/study-avs-may-not-detect-darker-skinned-pedestrians-as-often-as-lighter-ones/

[2] P. Koopman and M. Wagner, "Autonomous vehicle safety: An interdisciplinary challenge," *IEEE Intelligent Transportation Systems Magazine*, vol. 9, no. 1, pp. 90–96, 2017.

[3] R. McAllister, Y. Gal, A. Kendall, M. Van Der Wilk, A. Shah, R. Cipolla, and A. Weller, "Concrete problems for autonomous vehicle safety: Advantages of bayesian deep learning." International Joint Conferences on Artificial Intelligence, Inc., 2017.

[4] N. Kalra and S. M. Paddock, "Driving to safety: How many miles of driving would it take to demonstrate autonomous vehicle reliability?" *Transportation Research Part A: Policy and Practice*, vol. 94, pp. 182–193, 2016.

[5] E. Vinitsky, A. Kreidieh, L. Le Flem, N. Kheterpal, K. Jang, C. Wu, F. Wu, R. Liaw, E. Liang, and A. M. Bayen, "Benchmarks for reinforcement learning in mixed-autonomy traffic," in *Conference on Robot Learning*, 2018, pp. 399–409.

[6] C. Wu, A. M. Bayen, and A. Mehta, "Stabilizing traffic with autonomous vehicles," in *2018 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2018, pp. 1–7.

[7] H. T. Sullivan and M. T. Häkkinen, "Preparedness and warning systems for populations with special needs: Ensuring everyone gets the message (and knows what to do)," *Geotechnical and Geological Engineering*, vol. 29, no. 3, pp. 225–236, 2011.

[8] D. R. Montello and C. Sas, "Human factors of wayfinding in navigation," 2006.

[9] L. Head, S. Shladover, and A. Wilkey, "Multi-modal intelligent traffic signal system," *University of Arizona*, pp. 32–36, 2012.

[10] M. D. Simoni, K. M. Kockelman, K. M. Gurumurthy, and J. Bischoff, "Congestion pricing in a world of self-driving vehicles: An analysis of different strategies in alternative future scenarios," *Transportation Research Part C: Emerging Technologies*, vol. 98, pp. 167–185, 2019.

[11] S. Zang, M. Ding, D. Smith, P. Tyler, T. Rakotoarivelo, and M. A. Kaafar, "The impact of adverse weather conditions on autonomous vehicles: how rain, snow, fog, and hail affect the performance of a self-driving car," *IEEE vehicular technology magazine*, vol. 14, no. 2, pp. 103–111, 2019.

[12] P. E. Theberge, "Development of mathematical models to assess highway maintenance needs and establish rehabilitation threshold levels," *Transportation Research Record*, vol. 1109, pp. 27–35, 1987.

[13] M. Snaith and J. Burrow, "Priority assessment," *Transportation Research Record*, vol. 951, pp. 87–95, 1984.

[14] R. E. Asbahan, J. K. McCracken, and J. M. Vandenbossche, "Evaluating the cracking predicted by the mepdg using results from the sr 22 smart pavement study," 2008.

[15] P. Wang, S. Motamedi, T. Canas Bajo, X. Zhou, T. Zhang, D. Whitney, and C.-Y. Chan, "Safety implications of automated vehicles providing external communication to pedestrians," 2019.

[16] Q. Shen, "Location characteristics of inner-city neighborhoods and employment accessibility of low-wage workers," *Environment and planning B: Planning and Design*, vol. 25, no. 3, pp. 345–365, 1998.

[17] W. D. Montgomery, R. Mudge, E. L. Groshen, S. Helper, J. P. MacDuffie, and C. Carson, "America's workforce and the self-driving future: Realizing productivity gains and spurring economic growth," 2018.