White Paper: “Mathematical Challenges and Opportunities for Autonomous Vehicles” (Long Program, Fall 2020, Institute for Pure and Applied Mathematics)

Editors: Samitha Samaranayake (Cornell University), Benjamin Seibold (Temple University), Raphael Stern (University of Minnesota)

Authors: Montasir Abbas (Virginia Tech), Rahul Bhadani (Vanderbilt University), Monika Filipovska (University of Connecticut), George Gunter (Vanderbilt University), Ronan Keane (University of Michigan), Nour Khoudari (Temple University), Jana Kosecka (George Mason University), Hani Mahmassani (Northwestern University), Samitha Samaranayake (Cornell University), Benjamin Seibold (Temple University), Richard Sowers (University of Illinois at Urbana-Champaign), Jonathan Sprinkle (Vanderbilt University), Raphael Stern (University of Minnesota), Lawrence Udeigwe (Manhattan College), Daniel Work (Vanderbilt University), Jia Le Xian (Lizora Technologies), Mengsha Yao (Temple University)

Date of White Paper (after long program and two reunion workshops): December 1, 2023

Table of Contents:

1. Executive Summary

Autonomous vehicle (AV) research and development has achieved a similar status in terms of resources invested, societal excitement, and media coverage as space travel and exploration. At the same time, AV research is not rocket science; it is more complicated: while in itself, an AV is no more complex than a spacecraft, it must reliably interact and communicate with many other agents, particularly humans both inside and outside of the vehicle, much of it in a decentralized fashion. Hence, AVs, and their impact on us humans and our transportation systems, incur some of the most complicated science and engineering challenges that society will face in the near future. At the same time, there is some disconnect across the various research communities: professional product development is highly opaque, and public expectations and media communications have frequently been inaccurate or exaggerated.

This document summarizes outcomes, conclusions, and suggestions of the Long Program “Mathematical Challenges and Opportunities for Autonomous Vehicles”, organized by the Institute of Pure and Applied Mathematics (IPAM) from September 14 to December 18, 2020. Due to the COVID-19 pandemic, the long program was hosted fully virtually. This white paper also includes findings and discussions from two Reunion Workshops held in person at Lake Arrowhead in June 5-10, 2022 and June 11-16, 2023. During the second Reunion Workshop in June 2023, the participants discussed to summarize (1) key findings from the long program and reunion workshops, (2) what followup has occurred since the original long program, and (3) how this community can contribute to challenges in the development, deployment, and operation of autonomous vehicles in society.

We believe that this document can provide both communities, mathematics and AV-focused development, with an understanding of some of the key themes and challenges. Moreover, this document should be of interest to the broad collection of policymakers and analysts who will rely on qualitative tools to understand complex problems in the design, deployment, and operation of autonomous vehicles. Finally, the document can also be of interest to stakeholders of the national STEM effort in understanding how this community can fit into national needs and priorities, including equity in artificial intelligence (AI) and the application of AI in mission-critical roles.

A note on terminology: While there is a technical difference between autonomous vehicles and automated vehicles, in the context of this white paper, AVs are vehicles that drive on roadways and may have a high degree of automation (i.e., vehicles that are capable of driving without human intervention in a wide range of operational domains) or low degree of automation (such as driver-assist-enabled vehicles with ADAS features such as ACC).

2. Operation and Performance of Automated Vehicles

Introduction:

Cars in the future (even those designed to be autonomous) will likely build upon existing industry-wide vehicle architectures for data and control. Therefore, research in the area of vehicle automation must build on these industry standards and be compatible with accepted vehicle platforms. However, many research questions are still unanswered and there is room for improvement and innovation in developing reliable, safe, and efficient automated driving systems as outlined below. This section focuses on research needs specific to designing vehicle-level control to safely operate autonomous vehicles.

Subtopic: Modeling for autonomous driving

Much like traditional vehicles, automated vehicles rely on numerous individual components: all of the components on traditional vehicles, plus automation-enabling hardware like sensors, actuators, and compute infrastructure. While each individual component can be proved safe by combining models and tests, there is a gap in models and simulations that allow to test the full AV in which all components interact. To verify the true driving behavior of the AV and the safety of any controller, ideally before they can be implemented on physical vehicles, high-resolution models and simulations must be developed. In particular, there is a need for system design and testing approaches that use data from inter-process communication for training.

Subtopic: Deep learning in autonomous driving

Sensing and perception components that transform cameras into sensors are continuing to improve, but their errors need to be properly modeled to understand when there is a need to fuse with other sources and how to do that in real time. Additional sensing and perception modalities (radar, for example) also benefit from deep learning results through classification approaches.

Controllers are typically modeled through first principles. There are ongoing activities that involve learning-enabled control, which can use new modeled controllers to supervise learning controllers so they can be deployed at scale. This approach enables more explainable designs, which can improve the safety of vehicles as they are deployed.

A critical question going forward is how autonomous driving systems can make real-time safety-critical decisions based on complex sensing and perception data feeds (e.g., cameras, lidar, radar, etc.) by fusing both methods from artificial intelligence and traditional system models. This is a particularly challenging task since data can be unreliable, noisy, and reflect both the built environment and interactions with the consequences of human-made driving decisions on the road.

Subtopic: Co-design of AVs individually and in systems

The reliable design of future AVs will require insights into, and considerations of, how peripheral systems will be built. Examples include the electric charging infrastructure for future electric vehicles (EVs), which many expect to go hand-in-hand with AV development, infrastructure communication for vehicle-to-everything (V2X), as well as understanding how vehicles could interact with one another, purposefully through communication, as well as through emergent phenomena that are not designed. Moreover, many aspects of AV design and development reside in distinct research communities (e.g., mathematics, mechanical engineering, computer science, civil engineering, etc.), while the commercialization of AVs has distinct stakeholders (e.g., automotive manufacturers, municipalities, departments of transportation). Consequently, the existing work is too compartmentalized, and cross-silo collaboration, both in the research phase and in the implementation will be essential to design and deploy effective autonomous driving technology.

As with many emerging topics, research needs will evolve as new technologies become available, but understanding how technology will shape AV development, and vice versa, will be critical to design AVs that meet the needs of the consumer market and are practical for widespread adoption.

Subtopic: Safety

Verifiable safety guarantees are crucial for the broad societal adoption of AVs. Particular challenges that require fundamental advances include: perception of boundaries and obstacles, particularly in low-light environments using computer vision systems, inclusion of human intent in surrounding traffic, and the stability of real-time decision-making for AVs. These represent high technical challenges, and continued fundamental research is necessary to achieve satisfactory safety performance of automated driving systems.

Moreover, there is also a lack of consensus on what safety goals are socially acceptable. For instance, does it suffice if an AV drives as safe as a human driver, or does the bar need to be set higher than that? Or, if an automated system causes an accident, who is at fault and how are liabilities arranged? Or, should each AV be solely concerned with its own safe behavior, or should safety in an increasingly automated and connected environment become a centralized system-level objective? While AV safety is sometimes portrayed as a version of the ethical trolley problem, most challenges are in fact technical and legal.

3. Human Aspects of Automated Driving

Introduction: For the foreseeable future, humans will continue to be involved in the operation of automated vehicles. Therefore, there is a need to understand how humans will interact with, and use, automated driving systems. Some aspects that will need to be considered include how human drivers will drive in the proximity of AVs, as well as how the human operator will (1) interact with the vehicle, and (2) share control of the driving tasks with the AV. These aspects of human interactions with automated driving systems will need to be understood to enable safe and efficient deployment in an automated future.

Subtopic: Human interaction with automation

There is a need to understand how humans will interact with automated driving systems. This includes both understanding how human drivers will use AVs (as the ego vehicle), as well as how other road users will interact with such vehicles.

Since automated vehicles will likely remain lower-level (i.e., not capable of taking over all driving tasks) for the foreseeable near future, understanding how to effectively and safely switch control between the automated vehicle and the human driver (i.e., trust) needs to be fully understood.

Other drivers will also have to interact with automated driving systems. It will be necessary to understand how drivers will change their behavior if they know that vehicles are automated. Anecdotal evidence from cities with AVs operating on public roads have shown that in many cases, human drivers may take advantage of AVs, since they know these vehicles are not assertive and will avoid a collision. This motivates the need to study and mathematically quantify how human drivers will behave in the presence of other vehicles that are automated.

Finally, AVs will also have to interact with other non-motorized road users such as bicyclists, pedestrians, and wheelchair users. To safely operate AVs in such domains, accurate models are needed both for people's trust in AVs as well as models for AVs to estimate other road users' intent. For example, AVs may need to estimate a bicyclist's future trajectory, factoring in that the bicyclist may not signal their turn.

Subtopic: AVs and land use/urban design

Automated vehicles, particularly at higher levels of automation, may require dedicated infrastructure such as AV-only lanes, or AV-specific road signs. Understanding how human drivers will change their behavior with the addition of these infrastructure elements will need to be assessed.

Additionally, the need for downtown parking might decrease as AVs can drop passengers off and then drive themselves to remote parking facilities, or serve other riders. This could free up land for other uses such as create more walkable cities and neighborhoods. AVs may also reduce parking space required as they can be stored at a higher vehicle density. For example, human drivers need more space to open their doors after parking, while AVs do not. Designing urban and rural space allocation in an automated future will include questions on the optimal placement of AV parking facilities to minimize time to pick up passengers while leaving urban space for other land uses.

Finally, AVs may reduce the cost of travel for individual drivers, since they may be able to complete other tasks while driving. This may influence the future of the real-estate market since the demand for suburban living may increase as commuting becomes less burdensome. However, this may also increase vehicle miles traveled (particularly if many of these trips do not have human occupants in the vehicle), which may be a burden on local transportation agencies. Mathematical models for human adoption of AVs will be necessary to predict how the introduction of higher levels of autonomy will impact land use.

Subtopic: How to test AV safety with human agents

Human acceptance of AVs will depend on the perceived safety of automated driving systems. Therefore, there is a need for methods to quantify the safety of a specific driving action, as well as the need to estimate the safety of a particular automated driving system. Relying exclusively on roadway testing of such vehicles cannot suffice, because catastrophic events are inherently rare for individual vehicles on the road. Therefore, there is a need to develop mathematical tools to assess safety, as well as a need to design automated driving systems that provide better safety than a benchmark human driver.

When testing automated driving systems for safety, there is a need for data of humans interacting with automated vehicles. This includes vehicle trajectory data in a mixed-autonomy setting, as well as experimental data of drivers interacting with AVs as the AV operator. Previous efforts in the transportation community to collect large-scale datasets (e.g., Zen, I-24 MOTION, pNEUMA, TGSIM) have provided tremendous insight into traffic flow dynamics of traditional traffic/vehicles. Similar datasets that specifically include mixed-autonomy flow are needed to help capture interaction dynamics on local and macroscopic scales.

4. Technical Systems Interactions for Automated Driving

Introduction: The adoption of autonomous vehicles will not simply mean that humans that drive their cars manually will immediately vanish from the road. Rather, increasing vehicle automation and connectivity will fundamentally change the traffic patterns on our roads and also affect how safe our roads are. Thus, one major concern is ensuring that Connected and Automated Vehicles (CAV) will improve safety and enhance the overall traffic flow.

Subtopic: Understanding and predicting how connectivity and automation will influence traffic flow.

Mixed traffic streams remain a significant area of interest in terms of understanding and predicting how connectivity and automation will influence traffic flow. The major research areas can be divided into (1) modeling and simulation of vehicle agents and (2) changes in design and operation of the transportation network.

Modeling and simulation of vehicle agents: Modeling the interaction between human-driven vehicles and AVs requires not only modeling the AVs, but also modeling how humans interact with AVs compared to how they interact with other human-driven vehicles. Potential changes in human driving behavior may be a result of either attempting to take advantage of AVs in the traffic stream (e.g., saving travel time by forcing a lane change in front of AVs), or driving less efficiently due to excess caution when in the vicinity of AVs. Some humans might also intentionally attempt to disrupt autonomous driving due to the dislike of the technology.

Changes in design and operation of network: Design and operation of transportation networks in the presence of AVs will entail some additional aspects, including the consideration of the ranges of electric vehicles and the need for new energy infrastructure/resources. Placement of the new electrification and communication infrastructure needs to be optimized. In addition, the operation of the transportation network (e.g., managed lanes, traffic signals, ramp metering, etc.) will need to be optimized to consider the characteristics of AVs, and not only traditional vehicles' characteristics. There will be a room for new innovative designs and disruptive technology that can better address the new vehicle fleet and their handling characteristics. More research is needed to understand how simulation models should be adjusted to also account for electrification, and not just automation/connectivity.

Subtopic: Leveraging autonomy and connectivity to improve traffic operations

Automation and connectivity of vehicles can have significant desirable impacts on the roadway transportation system. Important examples include:

- AVs may be capable of following other vehicles at closer distances than a human can, resulting in significantly increased capacities (i.e., maximum admissible flow rates) of roadways.
- The conversion or re-dedication of existing HOV (high occupancy vehicles) or HOT (High occupancy toll) lanes to become dedicated AV-only lanes may allow the traffic in those lanes to realize significantly higher throughput and reduced delays than lanes with mixed autonomy or humans only.
- Truck platooning, i.e., trucks forming a "train" in which only the lead truck is traditionally driven, while the following trucks follow at close distances using automation, promises significant savings in operational costs like labor and fuel.
- AVs can be leveraged to dampen traffic waves and to smooth or harmonize the flow of traffic, resulting in reduced energy consumption and improved safety metrics (e.g., the CIRCLES project which was presented in some of the IPAM workshops).
- Coordinated merges, enabled via automation and connectivity, may reduce disruptions that traffic flow currently experiences caused by merging human-driven vehicles.
- Because AVs are themselves sensors, having more AVs on the roads will act as additional sources of real-time data that can be leveraged towards many forms of benefits.

At the same time, there are also key challenges and potential pitfalls if automation and/or connectivity is not implemented efficiently. For example, some AV/ACC systems have been demonstrated to potentially make traffic (more) unstable and make traffic waves more intense, resulting in less safety and higher energy demand.

Subtopic: Verifying that autonomous vehicles will operate safely in mixed autonomy settings

Humans are, for the most part, remarkably good at driving. On average, accidents occur at a rate of 1 per 1 million miles driven. Creating autonomous vehicles that can perform at this level of safety is a major technical challenge.

When evaluating safety, it is not sufficient to only consider the AV sensors and algorithms. Human behavior is also a key factor, because ultimately CAV will share the roads with human driven vehicles. The ultimate question we must consider is whether a world with additional automated driving will be more safe or less safe than the current world. Thus, a good evaluation of CAV safety should also consider effects such as different human behavior around CAV, and take into account both accidents that are the fault of CAVs, and accidents that are the fault of humans interacting with the CAV.

Besides the accident rate, we must also consider the crash severity, crash type, injury to road users, number of near-misses, and frequency of safety critical events. Additional metrics for quantifying safety can also be developed to better support research, regulation, and CAV development.

The safety for different road types and traffic situations should also be accounted for. If some driving situations are identified as less safe, more focus can be given to those situations. Examples of unsafe situations may include reacting to a lead vehicle suddenly decelerating, merging onto a busy freeway/highway, passing through an intersection when the traffic light changes to yellow/red, or instances where objects occlude CAV sensors. If cases where CAV operation might increase accident risk can be identified, then that information can guide decisions on when to engage/disengage the automation. Preferably, additional effort would be spent to improve safety performance in any identified high risk situations.

One critical point in the safe operation of CAV, is the issue of disengagement of automated driving systems. Current automation levels 3/4 systems, which by definition cannot operate in all situations, assume that the human occupant is 100% alert and ready to take over at a moment’s notice, but this is frankly entirely unrealistic. This “hand off” between the automation and the human is particularly concerning.

Another fundamental issue is that safety-critical events and accidents, are by definition, rare events in driving. Thus, those cases we care the most about, are also the cases we have the least data on. Integrated efforts involving mathematical models, traffic simulation, driving simulator studies, traffic datasets, and other resources are critical for studying these safety critical events. It is important to invest in further development and improvement of quality simulation tools that can be used for evaluating CAV safety performance without having to wait for accidents to endanger human drivers on real roads. Simulations are a vital tool here because they can heavily focus on the safety critical/crash events which occur only rarely in the real world. Virtual reality/mixed reality can be used as a tool to allow for real human input into simulations.

Overall, it is tempting to conclude that CAV, with their enhanced sensing capabilities, use of connectivity, and advanced compute capability, will be able to drive in a superhuman manner. However, research suggests that current AVs are actually less safe overall than human driven vehicles.

5. Societal Impacts of Automated Driving

While it is widely believed that roads with autonomous vehicles are likely to be the norm in a “reasonably near” future, the way we get to this future and how it evolves involves a number of tradeoffs and decisions. This discussion focuses on the need for the research, development and policy-making community to explicitly consider the societal impacts of a future with AVs. In particular, a key question is how we can quantitatively frame relevant discussions with a focus on the societal impact.

The general goal of transportation systems is to move people and goods at a societal scale with a focus on the potentially competing goals of safety, efficiency, accessibility, sustainability and reliability. The question arises as to whether AVs will provide us with an opportunity to improve the efficiency frontier with respect to these goals. Even if they were able to provide such an improvement, will we make the public policy, business and behavioral changes to reap these benefits? In the best case, AVs will improve road safety, increase our productivity, improve access to mobility, reduce environmental externalities etc. However, we recognize that such a future is not a given. It is also possible that we see a future where congestion is increased as more people commute from longer distances due to the reduced inconvenience of long commutes and increased cost of access to mobility further impact marginalized communities. While there has been a significant body of research on the technological underpinnings of AVs, the broader impacts of the technology have comparatively received less focus by the research community.

Subtopic: Lessons from ride-hailing

Ride-hailing services have provided a compelling example of how innovation can lead to private industry having a significant impact on mobility. The question of whether the impact has been positive for society in general is a complex one. Given this experience, how can this community help characterize the complex geospatial and economic tradeoffs needed to inform decision-makers who balance questions of incentives, regulation, and entrepreneurship, with the goal of facilitating the societal transition to autonomous mobility? One component of this is to understand how to frame questions and tradeoffs of social justice and equity in access to autonomous mobility.

Subtopic: Impact of AI

Effective autonomous mobility at scale will require a number of decisions to be made by AI. Questions of social justice and equity in training and assessing AI are becoming more and more important. Some of these questions are related to the growing body of work on fair AI. However, some of the questions that arise in AV systems have to do with fair allocation of resources and other problems that are not directly addressed by this work. Can this community help understand how these issues are manifested in autonomous mobility?

Concluding Remarks

Autonomous vehicles are on their way to revolutionize how we think about transportation systems. However, before they are ubiquitous, a wide range of technical challenges must be addressed and resolved as more and higher automated vehicles are being deployed on our roadways. These problems range from control of individual vehicles to how they interact with surrounding vehicles and pedestrians to how they meet the needs of larger geographic regions. Good mathematical models serving as “what-if machines”, simulations, measurements, and guarantees of performance are also needed in order for society and decision makers to think clearly about often-competing goals and objectives.

Automation of transportation, combined with its surrounding research needs, is a fantastic area for interdisciplinary research that spans the whole pipeline from mathematical foundations, over academic areas like engineering, computer science, and also social sciences, to industry, public stakeholders, etc. The IPAM long program, with participants from many of these fields, has shown that a productive interplay of all these areas is possible, and existing collaboration of participants have demonstrated that great outcomes can result from these interactions. The authors of this white paper would like to stress these significant opportunities for cross-disciplinary research with high broader and societal impact, and also to the opportunities for cross-disciplinary programs and initiatives, to funding agencies, public stakeholders, and education and research institutions. We hope that we can all help shape a successful transition to a better, safer, more efficient, more fair, and more enjoyable, future of transportation.

List of Acronyms

AV = Automated vehicle or autonomous vehicle (meaning dependent on context)

CAV = Connected and automated vehicle

STEM = Science, Technology, Engineering, and Math

V2X = Vehicle-to-everything

AI = Artificial intelligence

EV = Electric vehicle

ADAS = Advanced driver assistance systems

ACC = Adaptive cruise control