Optimizing Real-Time Autonomous Vehicle Control through Advanced Neural Networks:

Technical Refinements and Societal Implications

A Thesis Prospectus

In STS 4500

Presented to

The faculty of the

School of Engineering and Applied Science University of Virginia

In Partial Fulfillment of the Requirements for the Degree Bachelor of Science in Computer Science

By

Yuna Park

10/8/2023

On my honor as a university student, I have neither given nor received unauthorized aid on this assignment as defined by the Honor Guidelines for Thesis-Related Assignments.

ADVISORS

Kent Wayland, Department of Engineering and Society

Briana Morrison, Department of Computer Science

Introduction

Autonomous vehicles stand out as one of the most revolutionary emerging technologies, offering the potential to completely reshape transportation by enhancing safety, efficiency, and accessibility. However, developing reliable autonomous systems and seamlessly integrating them into existing transport infrastructure pose daunting challenges, both technical and societal.

This project aims to push the boundaries of real-time perception, planning, and control for autonomous vehicles through novel neural network architectures for multi-modal sensor fusion. By combining convolutional and recurrent neural networks, this research will explore optimized deep learning models to process spatial visual data along with temporal sequences from lidar and radar. These technical advancements are critical for public acceptance and safe integration of autonomous vehicles.

In addition, this project will critically evaluate the broader societal impacts of autonomous vehicles through an empirical STS analysis. Surveys, focus groups, and expert interviews will examine public perspectives on accountability and liability for accidents involving self-driving cars. With autonomous systems taking over more driving functions, questions around legal responsibility for crashes become increasingly complex. By investigating who people blame and how policies should adapt, this STS research can provide insights to guide the adoption of autonomous vehicle technologies.

Technical Research Problem - Unlocking Autonomous Vehicle Perception: Leveraging Multi-Modal Sensor Fusion and Deep Learning Architectures

Autonomous vehicles rely on accurate real-time perception of the environment to navigate safely. To enable robust spatial and temporal understanding, I propose leveraging deep neural networks for multi-modal sensor fusion across cameras, lidar, and radar. CNNs have demonstrated effectiveness in spatial perception tasks such as object detection and segmentation in camera data (Girshick et al., 2014; He et al., 2017). On the other hand, RNN architectures like Long Short-Term Memory (LSTM) and Gated Recurrent Unit (GRU) excel at modeling temporal sequences and have been successfully applied to lidar and radar data streams for prediction tasks (Hong et al., 2019; Kim et al., 2020).

CNNs, akin to intelligent eyes, are adept at processing grid-like data, particularly images and videos. They achieve object recognition capabilities by scanning images using small filters or kernels, which are designed to identify specific patterns like edges or textures. For example, one filter might recognize the curve of a cat's ear, enabling CNN to piece together the complete image, feature by feature. Kernels are essentially little matrix operations that traverse the input image to create activation maps, highlighting areas that match the pattern they are designed to detect.

On the other hand, RNNs are the memory wizards of the artificial intelligence realm. These networks excel at understanding sequential data, making them ideal for tasks involving time and order. Consider predicting the next word in a sentence. RNNs process each word sequentially, maintaining an internal memory of past inputs. For instance, given the sentence "The cat sat on the...," an RNN, using its contextual memory, might predict "mat" if it had previously seen "cat" or "cushion" if it had seen "dog." This sequential understanding allows RNNs to grasp the temporal aspects of data, making them invaluable for predicting movements over time.

In the realm of autonomous vehicles, CNNs serve as the vehicle's eyes, processing visual data from cameras to identify objects and their spatial locations. These networks enable the vehicle to 'see' the environment, recognizing pedestrians, vehicles, and other objects on the road. Simultaneously, RNNs analyze sequential data from sensors like lidar and radar. Lidar provides detailed 3D maps of the surroundings, while radar detects the speed and movement of objects. RNNs, especially variants like Long Short-Term Memory (LSTM) and Gated Recurrent Unit (GRU), excel at understanding this sequential data.

However, lack of diverse real-world multi-modal driving datasets makes training and evaluating such sensor fusion models difficult (Caesar et al., 2020). To help address this data scarcity, I will focus on studying leveraging simulation environments and data augmentation techniques. I will specifically study utilizing the hyper-realistic synthetic environments in CARLA and AirSim to generate diverse multi-modal driving data under various weather, lighting and traffic conditions difficult to capture through real-world driving alone.

CARLA (Car Learning to Act) is an open-source simulator built for development, training, and validation of autonomous driving systems (Dosovitskiy et al., 2017). It generates photo-realistic sensor data like camera images and lidar point clouds synced across a simulation environment modeled after real urban layouts. Weather, lighting and traffic conditions can be varied for data diversity. CARLA provides flexibility to extract paired data across modalities from many vehicle viewpoints.

AirSim (Shah et al., 2018) also offers high-fidelity visual and sensor simulations, while focusing more on support for aerial vehicles. It integrates with devices like Oculus for immersive data collection and provides programming interfaces in Python and C++. AirSim models complex environments like forests, mountains, and cities based on Unreal Engine, enabling rich lifelike scenario generation.

Models pre-trained on augmented simulation data may achieve better generalization to real-world driving (Müller et al., 2022). I will work on complementary environments by using CARLA for urban layouts and AirSim for aerial and sparse environments like mountains and forests. Using CARLA for street-level driving data and AirSim for aerial viewpoint data will provide variability in environments. Fusing the two can improve generalizability and the ability to extract perfectly paired camera, lidar and radar streams with precise ground truth labeling enables pre-training of sensor fusion models which can then be transferred successfully to actual autonomous vehicles. The models essentially leverage simulation as augmented training data before final deployment.

Throughout this research, I will be learning about how the combination of CNNs and RNNs addresses both spatial and temporal aspects of autonomous vehicle perception, and the effectiveness of data augmentation techniques in improving the diversity and richness of training datasets. I will also study how synthetic data generated in simulation environments can be augmented to cover a wide range of scenarios, including different weather, lighting, and traffic conditions. Finally, I will determine the best strategy for leveraging both CARLA and AirSim data, considering their strengths in simulating urban layouts and arial/sparse environments, respectively by evaluating how combining data from these environments improves the robustness and adaptability of the trained models.

STS Research Problem - Sociotechnical Implications of Autonomous Vehicle Adoption: Navigating Liability, Accountability, and Public Perception

The integration of autonomous vehicles into our transportation landscape holds transformative potential, promising improved mobility and efficiency. However, this promising future is accompanied by intricate sociotechnical challenges, particularly concerning liability, accountability, and public perception. This research will seek to explore and address these challenges to facilitate the societal adoption of autonomous vehicles.

Recent advancements in autonomous vehicle technology have sparked public skepticism, driven in part by deadly crashes that raise concerns about safety, ethical responsibility, and legal liability (Goodall, 2014). A 2021 survey by J.D. Power revealed that 67% of respondents expressed worries about liability rules for self-driving vehicles, highlighting a significant hurdle in public acceptance (J.D. Power, 2021). Such concerns are further fueled by critics who argue that accidents expose negligent practices prioritizing progress over safety (NTSB, 2019).

The central question guiding this research is: How should policies for autonomous vehicle liability align with public attitudes on accountability to facilitate widespread adoption? To unpack this complex issue, I will delve into the perspectives of various stakeholders, including the public, policymakers, autonomous vehicle manufacturers, legal experts, and technology ethics researchers.

A thorough examination of the existing literature reveals parallels with issues encountered in other domains, such as aviation and railroads In navigating the intricate landscape of autonomous vehicle adoption, a deeper exploration into the historical evolution of liability rules, as documented by The Rand Corporation in Smith's (2017) study, becomes paramount. This historical analysis offers a contextual foundation for comprehending how legal frameworks have adapted to the complexities introduced by technological advancements over time. Smith's work provides valuable insights into the challenges and responses in liability policies, creating a bridge between past and present considerations. By synthesizing this historical context, we can better grasp the nuances and dilemmas that autonomous vehicles present in terms of accountability.

Building upon this historical foundation, the study incorporates theoretical frameworks like robot ethics (Lin et al., 2012) and value-sensitive algorithm design (Friedman & Hendry, 2019). These frameworks offer lenses through which we can examine the ethical implications of autonomous systems and the intricate interplay between technology, responsibility, and societal values. Robot ethics, in particular, addresses the ethical considerations in the design and deployment of autonomous systems, emphasizing the need for responsible and accountable AI practices. Value-sensitive algorithm design extends this perspective by highlighting the importance of incorporating human values into the development of algorithms, especially those with societal impact. The debates on meaningful oversight policies for autonomous vehicles persist, with surveys indicating gaps between public attitudes toward accountability during crashes and limited corporate liability advocated by manufacturers (J.D. Power, 2021).

By intertwining historical analyses with contemporary ethical frameworks, the research aims to create a comprehensive narrative that not only explores liability policies but also delves into the ethical dimensions of autonomous vehicle technology. This integrated approach seeks to inform the central question of aligning policies with public attitudes on accountability. As we navigate these sociotechnical challenges, a holistic understanding that draws from both historical perspectives and contemporary ethical considerations becomes instrumental in shaping effective regulations and fostering public trust in the widespread adoption of autonomous vehicles.

This study will require a judicious selection of evidence to effectively address the central question of aligning policies with public attitudes on accountability in the adoption of autonomous vehicles. The primary sources of evidence will include an in-depth analysis of accident reports related to autonomous vehicles, an examination of liability rules in diverse operating domains, and targeted surveys and focus groups to capture public opinions and normative perspectives. Accident reports will be analyzed to discern the circumstances and backgrounds of incidents involving autonomous vehicles. This analysis will provide insights into the factors contributing to accidents and help identify patterns or commonalities that can inform liability policies. Focusing on relevant and representative case studies, this data will be instrumental in understanding the nuances of accountability in different scenarios. Additionally, an investigation into liability rules across various operating domains will contribute to the historical context of liability policies. By comparing existing policies in industries like aviation and railroads, the study aims to draw parallels and distinctions, identifying lessons learned and potential pitfalls in the evolution of liability frameworks. This historical analysis will serve as a foundation for shaping effective regulations in the context of autonomous vehicles. Surveys and focus groups will be conducted to gauge individual attitudes toward blame and accountability in autonomous vehicle accidents. Surveys will assess the degrees of culpability assigned to human and algorithmic components across various accident scenarios, providing quantitative insights into public perceptions. Focus groups, on the other hand, will offer qualitative data, capturing normative perspectives on accountable autonomy and shedding light on the societal expectations surrounding autonomous vehicle technology. To streamline the data collection process and ensure the manageability of the study, a focused approach will be taken. The emphasis will be placed on the most critical and informative sources: accident reports, liability rules, surveys, and focus groups. These sources collectively offer a comprehensive view of the sociotechnical landscape, enabling a nuanced exploration of liability and accountability in autonomous vehicle adoption.

By bridging the gap between policy and public perception, this research aims to contribute valuable insights to guide the development of effective regulations that align with societal expectations. Identifying areas for regulatory evolution can enhance public trust, thereby fostering increased readiness for the widespread adoption of autonomous vehicles.

Conclusion

In summary, this research prospectus outlines a comprehensive interdisciplinary investigation into the technical challenges of autonomous vehicles and their societal implications. The technical aspect focuses on enhancing neural network performance for real-time environmental perception using multi-modal sensor fusion. The research leverages simulation environments, such as CARLA and AirSim, to address the scarcity of diverse real-world driving datasets. Simultaneously, STS research component delves into public attitudes and liability policies. Through surveys, focus groups, and expert interviews, the study aims to inform policies for the responsible integration of autonomous vehicles. By bridging the gap between technology and society, this research strives to contribute novel AI techniques and empirical insights for the safe adoption of autonomous vehicles.

References

Caesar, H., Bankiti, V., Lang, A. H., Vora, S., Liong, V. E., Xu, Q., Krishnan, A., Pan, Y., Baldan, G., & Beijbom, O. (2020). nuscenes: A multimodal dataset for autonomous driving. 2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), 11621-11631. https://doi.org/10.1109/CVPR42600.2020.01162

Dosovitskiy, A., Beyer, L., Kolesnikov, A., Weissenborn, D., Zhai, X., Unterthiner, T., Dehghani, M., Minderer, M., Heigold, G., Gelly, S., Uszkoreit, J., & Houlsby, N. (2020). An image is worth 16x16 words: Transformers for image recognition at scale. arXiv preprint arXiv:2010.11929.

Friedman, B. & Hendry, D.G. (2019) Value Sensitive Design: Shaping Technology with Moral Imagination. New York: NYU Press.

Girshick, R., Donahue, J., Darrell, T., & Malik, J. (2014). Rich feature hierarchies for accurate object detection and semantic segmentation. 2014 IEEE Conference on Computer Vision and Pattern Recognition, 580-587. https://doi.org/10.1109/CVPR.2014.81

Goodall, N.J. (2014). Machine ethics and automated vehicles. In G. Meyer & S. Beiker (Eds.), Road vehicle automation (pp. 93-102). Springer.

Hong, S., Chang, T. H., Lee, H., & Shin, J. (2019). Classification of driving environments of autonomous land vehicle using lstm recurrent neural network. IEEE Access, 7, 70127–70134. https://doi.org/10.1109/ACCESS.2019.2919760

J.D. Power. (2021). Mobility Confidence Index Study. https://www.jdpower.com/business/press-releases/2021-us-mobility-confidence- index-study

Lin, P., Abney, K., & Bekey, G. A. (2012). Robot ethics: Mapping the issues for a mechanized world. Artificial Intelligence, 175(5-6), 942-949.

Müller, M., Savelli, F., Riedmaier, S., Schwesinger, U., & Knoll, A. (2022). Let Me Race Your Toy Car: Analyzing Real World LiDAR Data of a 1:10 Scale Racecar and a Full Size Autonomous Driving Car. 2022 IEEE Intelligent Vehicles Symposium (IV), 1288-1293.

NTSB. (2019). Collision Between Vehicle Controlled by Developmental Automated Driving System and Pedestrian, Tempe, Arizona, March 18, 2018 (Report No. HWY18MH010).

Shah, S., Dey, D., Lovett, C., & Kapoor, A. (2018). Airsim: High-fidelity visual and physical simulation for autonomous vehicles. In Field and service robotics (pp. 621-635). Springer, Cham.

Smith, B.W. (2017). Automated Driving and Product Liability. Michigan State Law Review 2017(1):1–80.