NeurIPS 2020 Workshop Proposal:
Machine Learning for Autonomous Driving

1. Workshop Overview

We propose a full-day workshop, called “Machine Learning for Autonomous Driving” (ML4AD), as a venue for machine learning (ML) researchers to discuss research problems concerning autonomous driving (AD). Our goal is to promote ML research, and its real-world impact, towards self-driving technologies. Full self-driving capability (“Level 5”) is far from solved and extremely complex, beyond the capability of any one institution or company, necessitating larger scale communication and collaboration, which we believe workshop formats help provide.

We propose a large-attendance talk format of approximately 500 attendees, including a call for papers with spotlight presentations and keynote presentations to communicate the current state-of-the-art; panel debates to discuss future research directions; a call for competition participation to encourage interaction around a common benchmark task; and social breaks for newer researchers within the autonomous driving (AD) community to network and meet others.

1.1 Importance of Autonomous Driving at NeurIPS

Autonomous vehicles provide a rich source of high-impact research problems for the machine learning community at NeurIPS in diverse fields including computer vision, probabilistic modeling, gesture recognition, pedestrian and vehicle forecasting, human-machine interaction, and multi-agent planning. The common goal of autonomous driving can catalyze discussion between these subfields, generating a cross-pollination of research ideas. Beyond the benefits to the research community, AD research can improve society by reducing road accidents; giving independence to those unable to drive; and inspiring younger generations towards careers in ML with tangible examples of ML-based technology clearly visible on local streets. During the COVID-19 pandemic, there is increased interest in sidewalk delivery robots also. As many NeurIPS attendees are key drivers behind AD-applied ML, this proposed workshop intends to bring researchers together from both academia and industries to discuss machine learning applications in autonomous driving.

1.2 Intellectual Excitement on the Autonomous Driving Topic

- **Academic excitement**: in 2019 NeurIPS featured at least 11 conference papers and 40 workshop papers on the topic of AD, a level of interest we expect to continue in 2020. However, AD is also a rapidly-changing field, and we expect AD research at NeurIPS 2020 to be filled with fresh new ideas compared to 2019. For example, perception has dominated AD research since the deep learning revolution (e.g. vision conferences ICCV / CVPR / ECCV have *multiple* AD workshops), but now we see increasing interest in tracking and forecasting topic (e.g. recent ArgoAI, Waymo, Interpret, Lyft competitions).

- **Industry excitement**: autonomous driving is now a multi billion dollar industry with all major automakers and rideshare companies investing in ML research. Given the complexity of the AD task is beyond any one company, large scale collaboration and dialogue are a necessity, leading to stronger industry interest in ML venues such as NeurIPS (e.g. Waymo, Uber, Lyft, Valeo, Wayve in 2019).
- **Previous engagement**: This ML4AD workshop proposal, if accepted, would be the 5th in a NeurIPS workshop series. Previous workshops were well attended by academia and industry in 2016, 2017, 2018, and 2019, the most recent receiving 45 submissions (40 accepted). AD is *increasingly popular* too: given the 2018 attendance of 300 people, ML4AD requested a 300 seat room in 2019, but was allocated a 550 seat room based on a 2019 workshop chair polling of NeurIPS registrants, which was mostly filled.
- **Public excitement**: AD is also constantly in the news, showcasing ML developments.

1.3 Challenge

To increase workshop engagement, we propose a new autonomous driving challenge to (1) stimulate new research ideas, (2) encourage researchers outside of NeurIPS to participate and present their methods, and (3) help advance the state-of-the-art research on a practical machine learning benchmark. The competition will be based on the CARLA Autonomous Driving leaderboard: https://leaderboard.carla.org, coordinated by German Ros and Valdlen Koltun, two of our workshop organizers and also authors of the CARLA simulator. Both German and Valdlen have prior experience coordinating AD challenges in 2019.

The main goal of this challenge is to evaluate the driving proficiency of autonomous agents in realistic traffic situations. Autonomous agents will have to drive through a set of predefined routes. For each route, agents will be initialized at a starting point and will be directed to drive to a destination point, provided with a description of the route. Routes will happen in a variety of areas, including freeways, urban scenes, and residential districts. In each of these routes, agents will have to deal with challenging traffic situations based on the NHTSA pre-crash typology. On top of this, agents will have to cope with a variety of weather conditions, including daylight scenes, sunset, rain, fog, and night, among others.

This competition will be deployed online by using Amazon AWS infrastructure. Teams will sign up in the platform and send their solutions as docker containers. Each docker container will be automatically evaluated by the CARLA leaderboard according to a set of driving metrics. Participants will be able to compete in two categories:

- SENSORS: only cameras, LIDAR, RADARs, IMU, and GNSS sensors are allowed.
- MAP: in addition to the above sensors, HD maps are available.

Participants will be able to join the CARLA Autonomous Driving community, which provides a space to exchange ideas and ask for help through a Discord platform [link] and a Discourse forum [link]. These spaces will be used to promote the engagement between the participating teams and the organizers.

2. Confirmed Speakers

The following nine speakers have each confirmed that they are able to present. This includes speakers that are also from outside the NeurIPS community (symbolized with a * symbol):

Name	Institution	Country	Position	Website	Research
Angela Schoellig	UToronto	Canada	Assistant Professor	website	scholar
Dragomir Anguelov	Waymo	US	Principal Scientist	website	scholar
Beipeng Mu *	Momenta.ai	China	R&D Director	website	scholar
Ehud Sharlin *	UCalgary	Canada	Professor	website	scholar
Pin Wang *	UCBerkeley	US	Research Associate	website	scholar
Jianxiong Xiao	AutoX	US,China	CEO	website	scholar
Sertac Karaman	MIT,OptimusRide	US	Associate Professor	website	scholar
Patrick Perez	Valeo	France	Scientific Director	website	scholar
Byron Boots	UWashington	US	Associate Professor	website	scholar

3. Diversity and Inclusion

We specifically selected speakers with a diverse technical **viewpoints** of AVs in both industry and academia on: perception (Jianxiong Xia, Patrick Perez); forecasting (Dragomir Anguelov); mapping (Beipeng Mu), planning (Pin Wang, Angela Schoellig), control (Byron Boots), explainability and human-machine interaction (Ehud Sharlin). In addition, we aimed to fairly represent the demographics within the NeurIPS community including: **gender** (3 women and 6 men); **race**; **geography** (Canada, China, France, US); **affiliations** (5 academia, 4 industry); and **seniority** (researcher, assistant professor, associate professor, professor). Our organizers are also diverse in gender, race, and seniority.

3.1. Encouraging Discussion

Beside keynote talks, we plan to encourage broad discussion in four ways:

- **Spotlight talks**: since NeurIPS will be virtual this year, we will allocate *every* accepted paper a “spotlight talk” to give exposure to all authors’ work.
- **Panel discussion**: to encourage lively debates, controversial issues, better ways of thinking about current problems, and also problems that are important but currently under-explored as suggested research directions for newer researchers to consider
- **Competition**: to actively stimulate new ideas, and encourage newcomers to NeurIPS, with discussion facilitated by a common benchmark in which to compare methods and discuss differences with multiple entrants invited to present their solution.

    - **Lowering the bar to entry**: even in simulation, writing a complete AD system can take months. We therefore aim to provide "starter template code" using either our Learning by Cheating code, or the PyLot open source software.
- **Social**: via Online Town or Gather Town (2D virtual environments) for people to meet.

3.2. Ease of Access

- **Streaming access**: many online services are not available in all nations. For example, YouTube Live streaming is not available in China. Depending on the virtual platform NeurIPs uses this year, we plan to use additional streaming services to ensure anyone in the world can view the workshop, such as https://www.bilibili.com or https://meeting.tencent.com which are available to Chinese attendees.
- **Website access**: we will host a website at https://ml4ad.github.io/ (updating it for 2020 if accepted) since github is accessible in every nation including China. Our website will list the talk titles prior to the event, clearly state there will be no archival proceedings for submissions, and list contact emails and a FAQ as 2019 did.

4. Timeline

- Call for Papers: Mon, Aug 3rd
- Submission: Wed, Sept 16th
- Notification: Wed, Sept 30th
- Camera Ready: Wed, Nov 25th
- Workshop: Fri Dec 11th, or Sat 12th

5. Schedule

A tentative schedule in Vancouver-time is:

Morning Times	Event	Afternoon Times	Event
08:50	Welcome	1:00	Keynote 5
09:00	Keynote 1	1:30	Keynote 6
09:30	Keynote 2	2:00	Spotlight Talks
10:00	Spotlight Talks	2:40	Coffee break
10:40	Coffee break	3:00	Keynote 7
11:00	Keynote 3	3:30	Keynote 8
11:30	Keynote 4	4:00	Keynote 9
12:00	Lunch	4:30	Panel

6. Organizers

We believe our organizational team below is well-suited to conduct this workshop, given both the diversity of our disciplines as well as a common interest grounded in autonomous driving applications. Our team comes from machine learning, computer vision, and simulation systems. We have experience organizing successful workshops in the past, discussed below.

Rowan McAllister rmcallister@berkeley.edu
is a postdoctoral researcher at UC Berkeley working on motion planning for autonomous vehicles. He previously worked with Uber's autonomous vehicle forecasting team, completed his PhD at the University of Cambridge on Bayesian reinforcement learning, and a masters in motion planning at the Australian Center for Field Robotics. Workshop organizing experience:

- NeurIPS 2019 workshop on Machine Learning for Autonomous Driving
- ICLR 2019 workshop on Task-Agnostic Reinforcement Learning
- RSS 2020 workshop on Interaction and Decision-Making in Autonomous-Driving
- ICML 2020 workshop on AI for Autonomous Driving
- ECCV 2020 workshop on Perception for Autonomous Driving

Nick Rhinehart nrhinehart@berkeley.edu is a postdoctoral researcher at UC Berkeley. Nick received his PhD in Robotics from Carnegie Mellon University. His work focuses on fundamental and applied research in machine learning and computer vision for behavioral forecasting and control in complex environments, with an emphasis on imitation learning, reinforcement learning, and deep learning methods.
Workshop organizing experience includes:

- ICML 2019 workshop on Imitation, Intent, and Interaction
- NeurIPS 2019 workshop on Machine Learning for Autonomous Driving
- CVPR 2018 tutorial on Inverse RL for Computer Vision

Xinshuo Weng xinshuow@cs.cmu.edu
is a PhD student at Carnegie Mellon University working on 3D Computer Vision and Graph Neural Networks in the context of Autonomous Driving. She previously worked at Facebook Reality Lab as a research engineer on high-fidelity full body human reconstruction, and completed her master degree at Carnegie Mellon University on Object Detection and Tracking.

Daniel Omeiza daniel.omeiza@cs.ox.ac.uk
is a PhD student at the University of Oxford working on explainability in autonomous vehicles. He is also a research candidate in the mobile robotics group of the Oxford Robotics Institute. He obtained a master's degree from Carnegie Mellon University and has worked for IBM Research as a research intern on AI explainability. He served as a volunteer in the Black in AI (BAI) workshops co-located with the NeurIPS conference in 2018 and 2019.

Fisher Yu i@yf.io
has a position of tenure-track assistant professorship in computer vision at ETH Zurich. He pursued his Ph.D. degree at Princeton University, advised by Thomas Funkhouser. His research interest lies in representation learning for image recognition, internet-scale visual understanding, interactive data processing system, and high-level understanding of dynamic 3D scenes. Workshop organizing experience includes:

- CVPR 2017, 2018, 2019, 2020 workshop on Autonomous Driving
- NeurIPS 2019 workshop on Machine Learning for Autonomous Driving
- ICML 2019 workshop on Human In the Loop Learning
- NeurIPS 2016 workshop on 3D Deep Learning
- CVPR 2015, 2016 workshop on Large-scale Scene Understanding Challenge

German Ros german.ros@intel.com
is a research Scientist at Intel Intelligent Systems Lab, working on Machine Learning, Simulation, Virtual worlds, Transfer Learning, and Intelligent Autonomous agents. He leads the CARLA organization and manages the Open3D project. Before joining Intel Labs, German served as a Research Scientist at the Toyota Research Institute, where he conducted research in the area of Simulation for Autonomous Driving, Scene Understanding, and Domain Adaptation, in the context of Autonomous Driving. Workshop organizing experience includes:

- ECCV 2016 Virtual/Augmented Reality for Visual Artificial Intelligence
- CVPR 2019 CARLA Autonomous Challenge workshop
- SIGGRAPH 2019 workshop on Computer Graphics for Autonomous Vehicles
- CVPR 2020 workshop on Embodied AI

Vladlen Koltun vladlen.koltun@intel.com
is the Chief Scientist for Intelligent Systems at Intel. He directs the Intelligent Systems Lab, which conducts high-impact basic research in computer vision, machine learning, robotics, and related areas. He has mentored more than 50 PhD students, postdocs, research scientists, and PhD student interns, many of whom are now successful research leaders.

7. Key Differences of this Workshop

7.1. Differences the 2019 version of the workshop:

Our organizational team for 2020 includes **four new organizers** compared to the NeurIPS 2019 workshop on Machine Learning for Autonomous Driving to encourage fresh new ideas. Specifically, **Daniel Omeiza** has promoted the topic of *explainability* of AV decisions, leading to securing Ehud Sharlin as a speaker of AV explainability, a professor who normally publishes outside of NeurIPS. Both **German Ros** and **Vladlen Koltun** are helping coordinate a new CARLA challenge integrated into the workshop. In 2019 offline perception and prediction challenges, but this CARLA challenge tests online control algorithms, which we hope will increase audience engagement. **Xinshuo Weng** is also new and will help coordinate our review process.

7.2. Differences with other workshops:

NeurIPS has many workshops, but this ML4AD workshop series has a unique focus on ML for transport and vehicles during 2016, 2017, 2018 (under the previous name of “Intelligent Transportation Systems”), and 2019.

8. Implementation

8.1. Managing conflicts of interest

We will avoid conflicts of interest when assessing submitted contributions by ensuring no reviewer conflicts with their assigned papers in terms of overlap of their recent institutions using the CMTautomatic functionality to check for conflicts. Last year, the program committee comprised 51 researchers, making it easy to find non-conflicted reviewers for each submission. As organizers, we will base our acceptance decisions on the committee’s recommendations, and not present any of our own works. To reduce variance in reviews, we will make the reviewer questions available online in advance to authors.

8.2. Soliciting Participation

If accepted, we will encourage participation by advertising the call for papers and competition through multiple channels, including public mailing lists (ml-news, robotics-worldwide), social media (Twitter, LinkedIn), internal mailing lists from current and previous institutions, personal websites, and contacting colleagues and researchers in the field inviting them to participate.