Project ID# eems062
Pillar(s): CAVs

Deep-Learning for Connected and Automated Vehicle (CAV) Development

Robert M. Patton, Peter Graf

Oak Ridge National Laboratory
National Renewable Energy Laboratory

2020 DOE Vehicle Technologies Office

Annual Merit Review Presentation

June 2, 2020

Overview

Timeline

- Project start: October 1, 2018
- Project end: September 30, 2021
- Percent complete: 50%

Budget

- Total project funding
  - DOE share: 100%
  - Contractor share: 0%
- Funding received in FY 2019
  - $ 1.9M
- Funding received in FY 2020
  - $ 2.2 M

Barriers

- Proprietary or expensive simulation tools and imagery for development of autonomous vehicle algorithms
- Significant expertise and manual effort required to develop machine learning algorithms
- Machine learning algorithms developed using desktop computing power

Partners

- Lead: Oak Ridge National Laboratory (ORNL)
- National Renewable Energy Laboratory (NREL)
- General Motors (GM)

Relevance

- Challenges
  - Much research in machine learning for CAVs is heavily focused on sensing / perception, and is often isolated from other aspects such as control or communication
  - Machine learning for CAV operation was initially heavily focused on safely operating according to traffic control structures first formed in the early 1900s with little to no concern for energy efficiency
  - Further exploration of machine learning for energy efficient CAV operation is needed
- Objective
  - Demonstrate HPC-based ability to analyze large data sets from prototype self-driving vehicles and discover higher performance and resilient operating algorithms for sensing, perceiving, and control.
  - Develop and demonstrate new machine learning based algorithms for vehicle operating controls that are capable of scaling to “Level 5” autonomous vehicle capabilities.
  - Develop a virtual test environment capable of training & safely evaluating autonomous vehicle operating controls over millions of miles and scenarios/environments expected to be encountered
  - Demonstrate workflow combining Imitation Learning with Reinforcement Learning

Milestones

Milestones	FY20 Q1	FY20 Q2	FY20 Q3	FY20 Q4
1. Utilize at least 10% of an HPC system resource simultaneously in a single CAVs learning task				
2. Improved perception performance on an existing open dataset.				
3. Develop new approach for training machine learning algorithms in a virtual environment				
4. Demonstration of ML-based coordinated autonomous vehicles control strategy (e.g. driving through intersection, curb-side pickup, intelligent merging); Generation of preliminary synthetic data set that would support multiple machine learning techniques.				

Any proposed future work is subject to change based on funding levels

This presentation does not contain any proprietary, confidential, or otherwise restricted information

Approach (1)

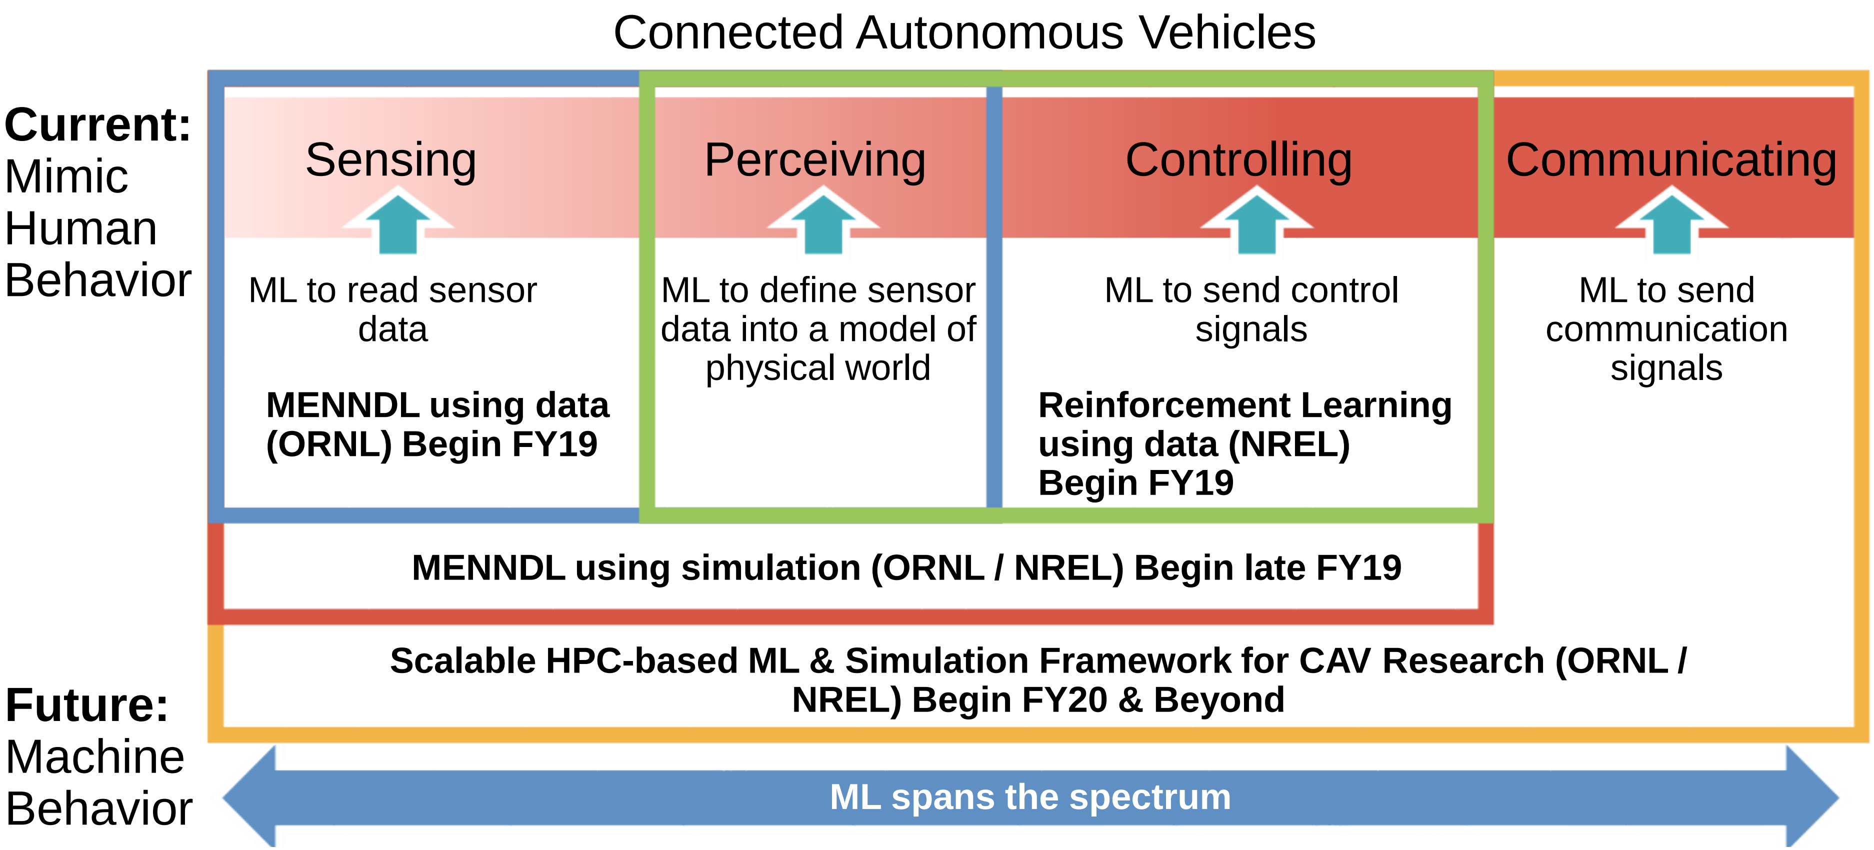

Goal

Develop a computational capability that leverages **Modeling & Simulation, High Performance Computing, and Artificial Intelligence** in order to enable the rapid development of perception, control, and communication algorithms for CAV

Technical Accomplishments and Progress - Summary

- MENNDL for Object-Perception
- Imitation Learning for Perception & Control
- Training Data for Imitation Learning
- Adversarial Testing & Evaluation
- Bridging 3D Driving Simulation to 2D Traffic Simulation - KRoad
- KFlow: Integration of Physics Based Driving & Reinforcement Learning in SUMO
- Transfer Learning: from 2D to 3D
- Integration of Imitation Learning with Reinforcement Learning

MENNDL for Object Perception

- MENNDL compared to baseline network called HRNet
  - J. Wang et al., "Deep High-Resolution Representation Learning for Visual Recognition," in IEEE Transactions on Pattern Analysis and Machine Intelligence. doi: 10.1109/TPAMI.2020.2983686
- Evaluation being performed using CityScapes data and data provided from GM that has similarities to CityScapes
  - https://www.cityscapes-dataset.com/
- Initial results show comparable performance but MENNDL approach is more tunable

	CityScapes	GM Data
HRNet (except no pretraining)	0.05-0.12	0.30
MENNDL (no pretraining)	0.14	0.29
HRNet (with pretraining)	0.38	0.51
MENNDL (with pretraining)	TBD	TBD

Imitation Learning for Perception & Control

- Using conditional imitation learning to perform end-to-end driving
- Behavior can be improved by altering the training data alone
- Driving behavior is assessed based on quantitative metrics
  - E.g., measuring distance from center of lane and distance traveled

Town 1 in CARLA

Town 3 in CARLA

Training Data for Imitation Learning

- Training data for imitation learning plays a very significant role in creating appropriate driving behavior
- Randomly generated data has implicit bias toward driving straight
  - Initial data set consisted of 97% straight with 3% consisting of turns
- Training data needs to be “engineered” for machine learning
  - e.g., more data from turns or underrepresented scenarios
  - Increasing distribution of turns up to 10% significantly improves driving behavior
- Research question: What training data distribution is needed in order to generalize to “any” city?

Gremlin: Adversarial Testing & Evaluation

- Find the scenarios that cause the autonomous vehicle to fail
  - Weather & lighting conditions
  - Vehicle & pedestrian traffic conditions
  - Various city & road layouts
- Leveraging CARLA & ORNL's Summit
- Generate millions of simulated scenarios & miles driven
- 2020 – 2021 DOE ALCC Proposal: *Evolutionary Multi-scenario Simulation Environment for Autonomous Vehicle Testing*
  - Partnered with GM; proposed 250K compute hours on ORNL's Summit

Bridging 3D Driving Simulation to 2D Traffic Simulation - KRoad

- When exploring AI opportunities across spectrum of CAVs control tasks, use the fastest simulator (simplest model) necessary for task
- Not all tasks in the CAVs pipeline require high fidelity visual 3D simulators such as CARLA
- The simpler models are a “bridge” enabling interoperability of AI algorithms and learned agents between 3D CARLA and 2D traffic simulators such as SUMO
- Call this “bridge” - KRoad

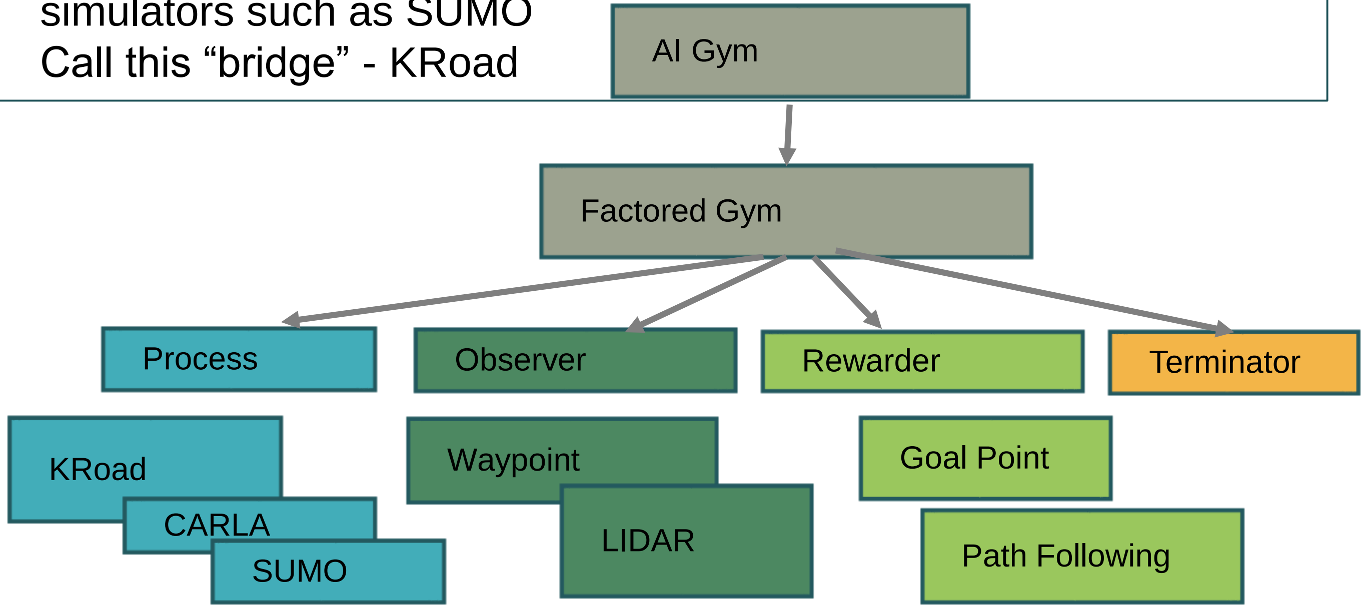

We have implemented the 2D dynamic model (DBM) within a flexible python framework (the “FactoredGym”) allowing for a mixing and matching of driving scenarios, physical models, and learning algorithms.

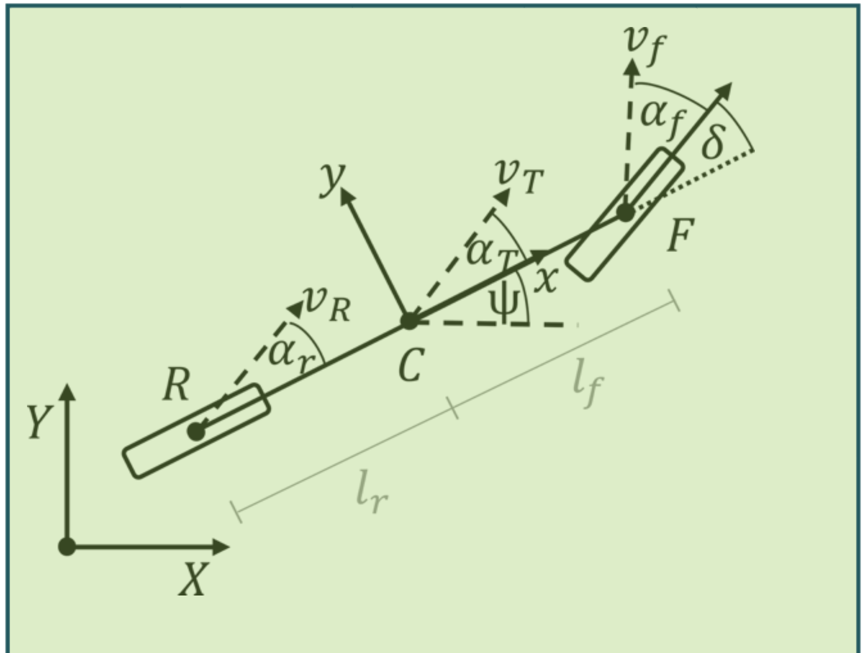

Geometric model of DBM

Using an industry/research standard system of ordinary differential equations governing the 2D physics of a vehicle with the above representation, we evolve these equations forward in time with steering and acceleration commands provided by our FactoredGym RL agent, or by a classical controller

Path Tracking in KRoad

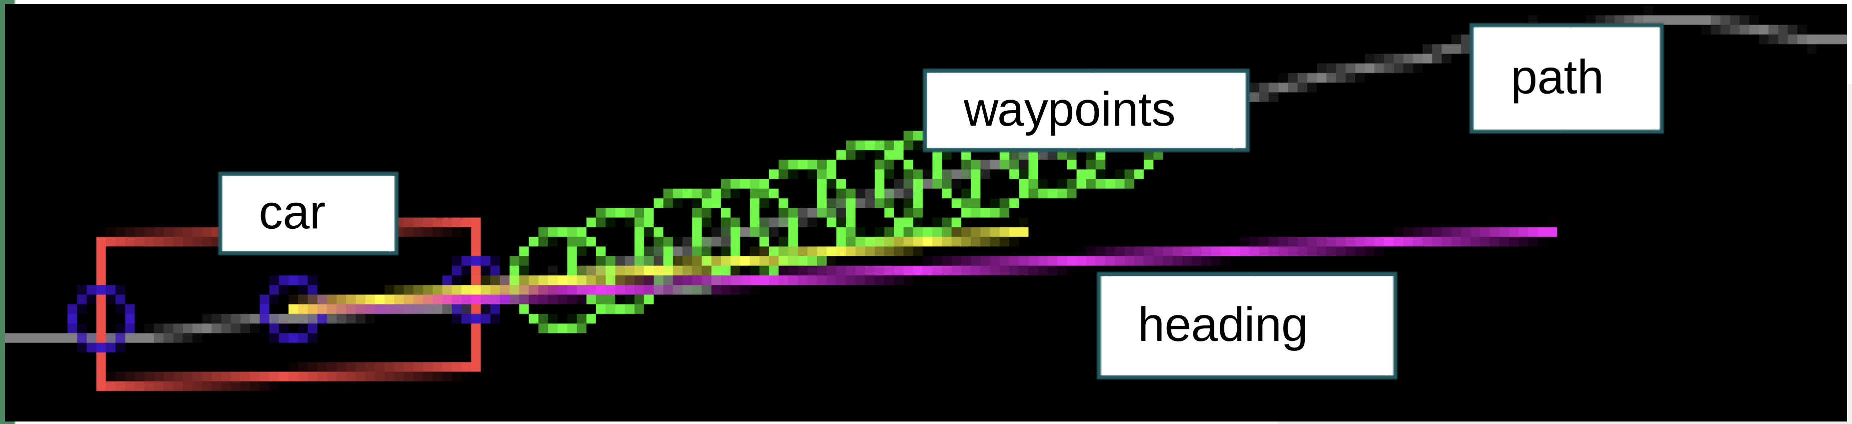

Trained Reinforcement Learning agent navigates random path

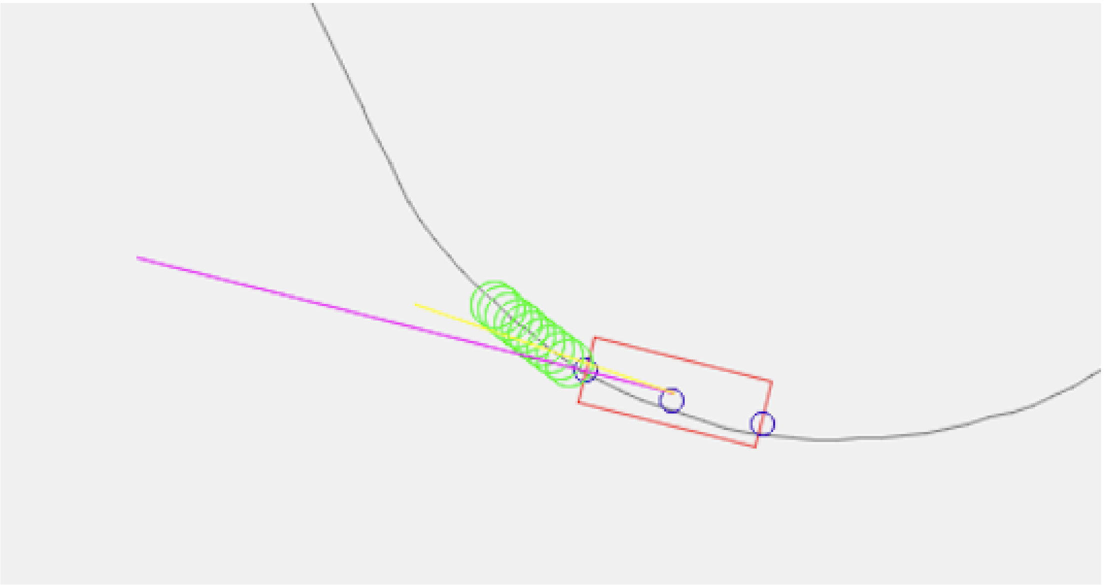

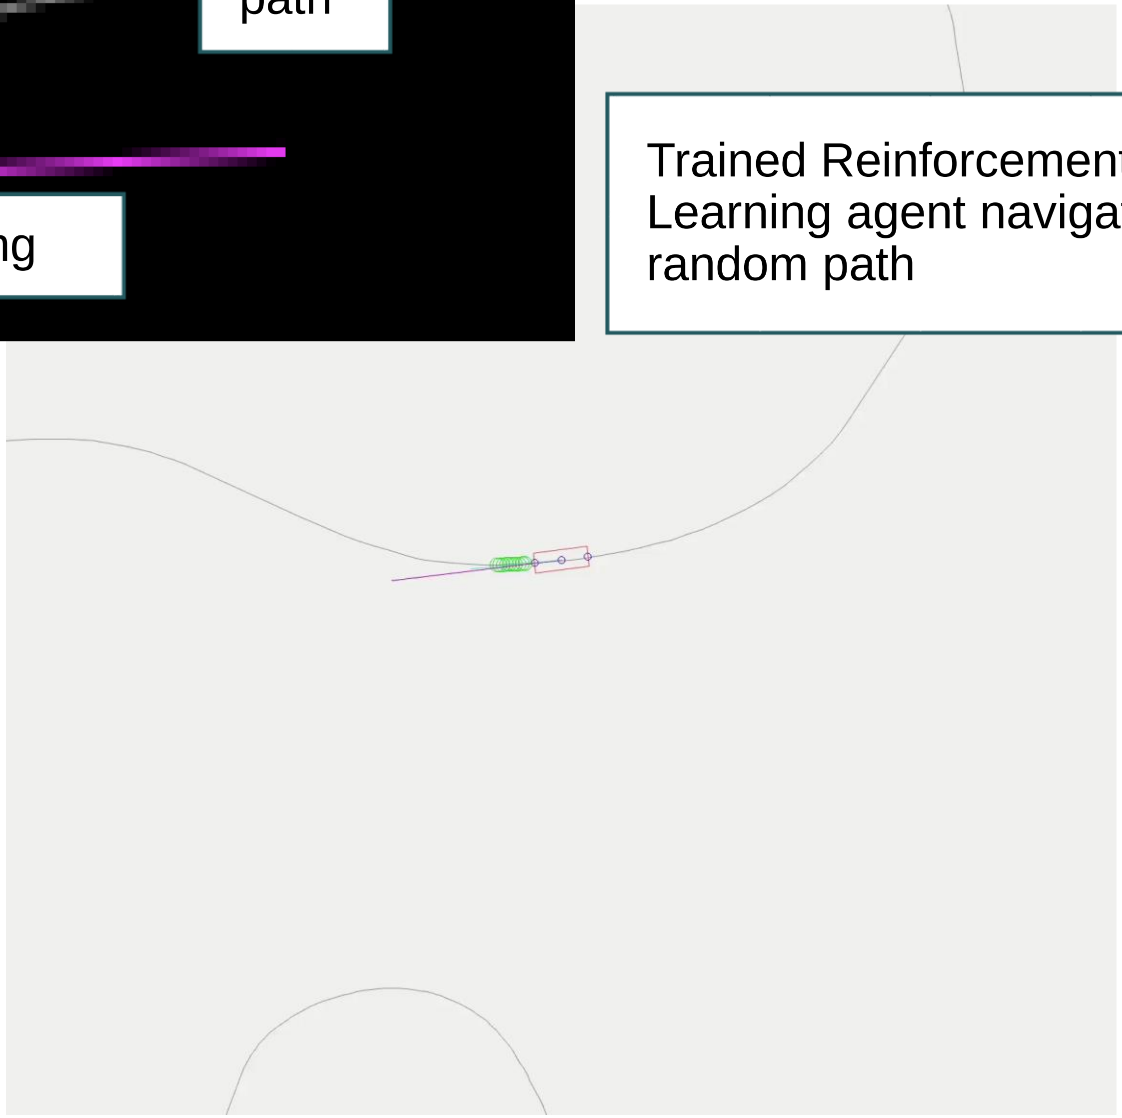

KFlow: Integration of Physics Based Driving & Reinforcement Learning in SUMO

- SUMO uses unrealistic homogenous driving algorithms for each car
- KFlow enables more realistic driving algorithms to be trained and used within SUMO (e.g., merging into a lane as opposed to simply “jumping” to a new lane)
    - Based on Flow, which enables reinforcement learning in SUMO

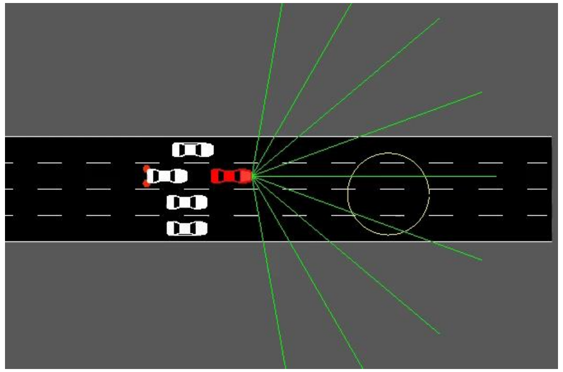

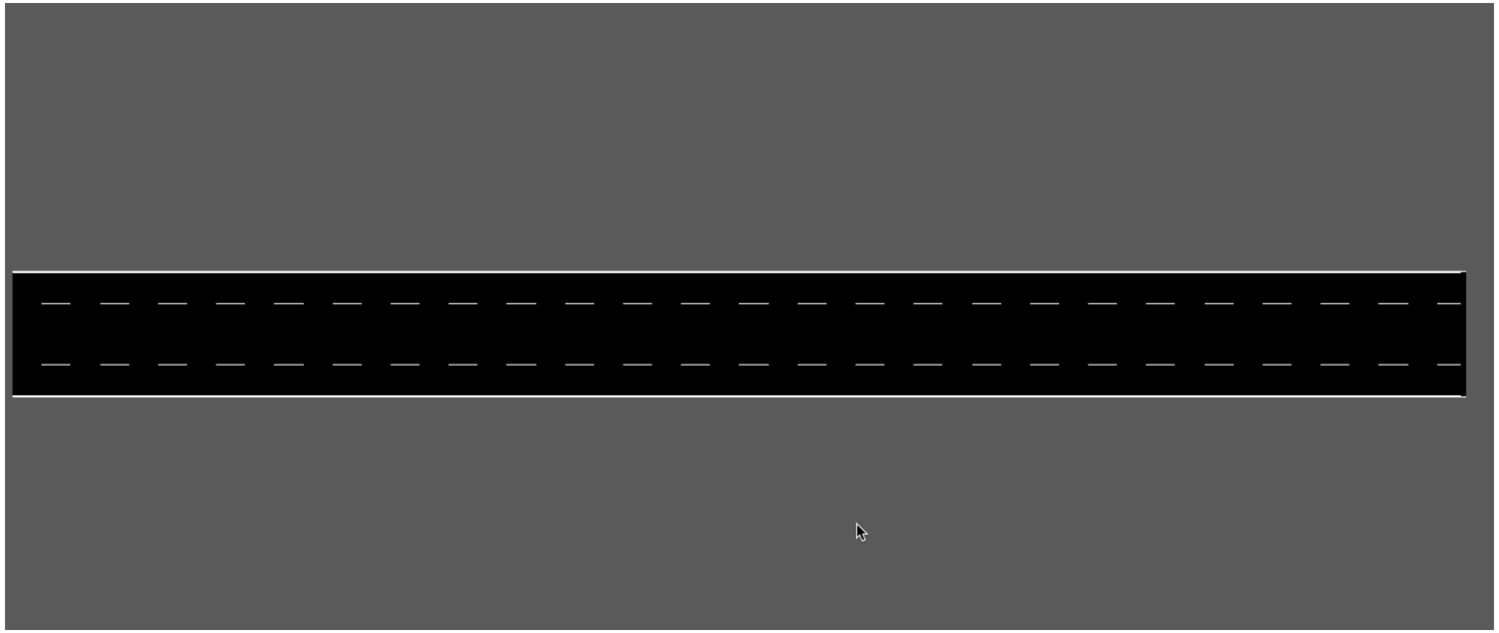

- SUMO gym environment that recreates the necessary elements for RL using continuous acceleration and steering in a 2D physics-based model.
- The vehicle is being directed to drive to the circle

Vehicle is capable of learning and environment provides important observations such as other vehicle speeds and these lidars that show in green that identify other lanes

Transfer Learning: from 2D to 3D

We are integrating KRoad and CARLA such that CARLA maps can be exported to KRoad, where 2D learning can occur in a fast environment, and the resulting controller then operated in 3D in CARLA. Eventually these can become pretrained models that are refined in CARLA based on more complex inputs (e.g. sensor data rather than waypoints).

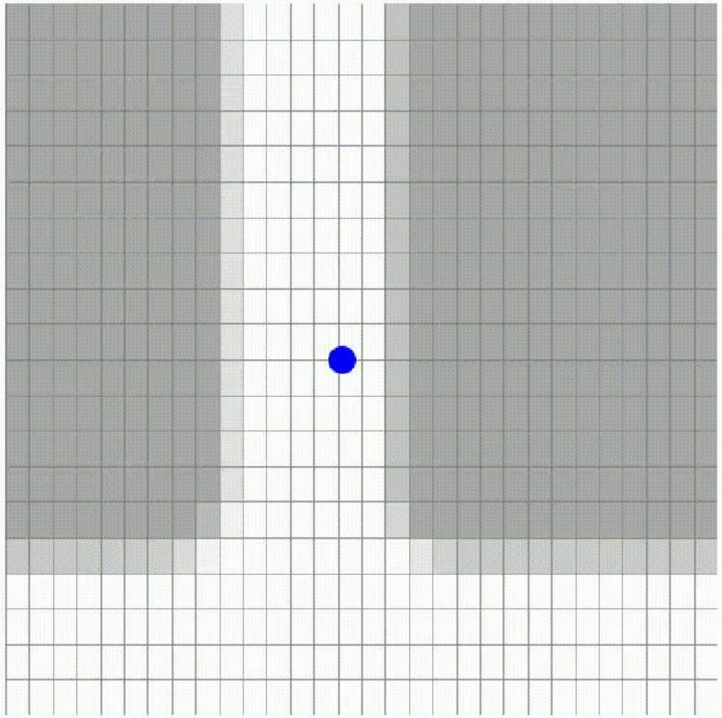

Responses to Previous Years Reviewers Comments

- Reviewer comments: need an industry partner
  - Response: We have partnered with General Motors. We have bi-weekly phone calls with them to discuss our progress and to brainstorm solutions to technical challenges. We also work with the CARLA development team, which is funded by Intel, Toyota, and General Motors.
- Reviewer comments: leverage existing realistic data sets
  - Response: Based on our interactions with General Motors, we are using the CityScapes data set as well as a data set provided by GM, which has characteristics very similar to CityScapes.
- Reviewer comments: need a clear benefit to energy usage reduction
  - Response: We are working toward demonstrating energy reduction. However, to do so, we must first provide a fundamental understanding of how machines can be effectively trained to drive a vehicle before they can be trained to reduce energy. Our long-term goal is to reduce overall fuel consumption through consistent driving behavior.
- Reviewer comment: “Using an open source software that might be very far behind commercial packages may be slowing the project down. Also, with as active as this area is, the reviewer would expect that companies would be lining up to be part of this work.”
  - Response: Commercial packages are extremely expensive (e.g., 5 to 6 figures for a single license per year) and typically only run on a desktop computer. They are not built or intended for high performance computing. We are using CARLA, not only because it is open source but because it is also funded by Toyota and GM. Companies are interested in what they can take from the project, and not what they can provide. GM has been an exception to that experience. We are witnessing the birth of a new multi-billion dollar industry, so most companies are not sharing any useful information beyond pure marketing content.

Remaining Challenges and Barriers

- Challenges and Barriers to both the project and the field:
  - Existing software codes currently built with “desktop” compute power in mind
  - Machine learning algorithms require significantly more compute time than anticipated
  - CARLA development controlled by outside entities & actively being developed with new features and bugs being released very quickly

Proposed Future Research FY21

- **Scalability**:
  - How many different scenarios (*not* miles driven!) can be achieved?
  - How many different driving algorithms can be evaluated?
- **Trainability**:
  - How to reduce time to solution for training?
  - How to reduce volume of training data or simulation time?
- **Desired Machine Behaviors**:
  - How do machines drive for increased energy efficiency?
  - How do machines drive in severe weather conditions (e.g., blizzard, high winds)?

Any proposed future work is subject to change based on funding levels.

Summary Slide

- **Approach:**
  - Develop a computational capability that leverages Modeling & Simulation, High Performance Computing, and Artificial Intelligence in order to enable the rapid development of perception, control, and communication algorithms for CAV
- **Technical Accomplishments (slide 6):**
  - MENNDL for Object-Perception
  - Imitation Learning for Perception & Control
  - Training Data for Imitation Learning
  - Adversarial Testing & Evaluation
  - KRoad
  - Kflow
  - Transfer Learning: from 2D to 3D
- **Future Work FY21:**
  - Enhanced perception and control algorithms using MENNDL using supervised and imitation learning
  - Scaling and integration of tools supporting Reinforcement Learning for driving simulators
  - Scaling up the computing resources and scenario generation

QUESTIONS?

Technical Backup

Integration of Imitation Learning with Reinforcement Learning

- Leveraging state of the art machine learning research in gaming, we have begun development of using a neural network trained with Imitation Learning first followed by Reinforcement Learning
  - Much like a person learning from someone else and then trying the task on their own

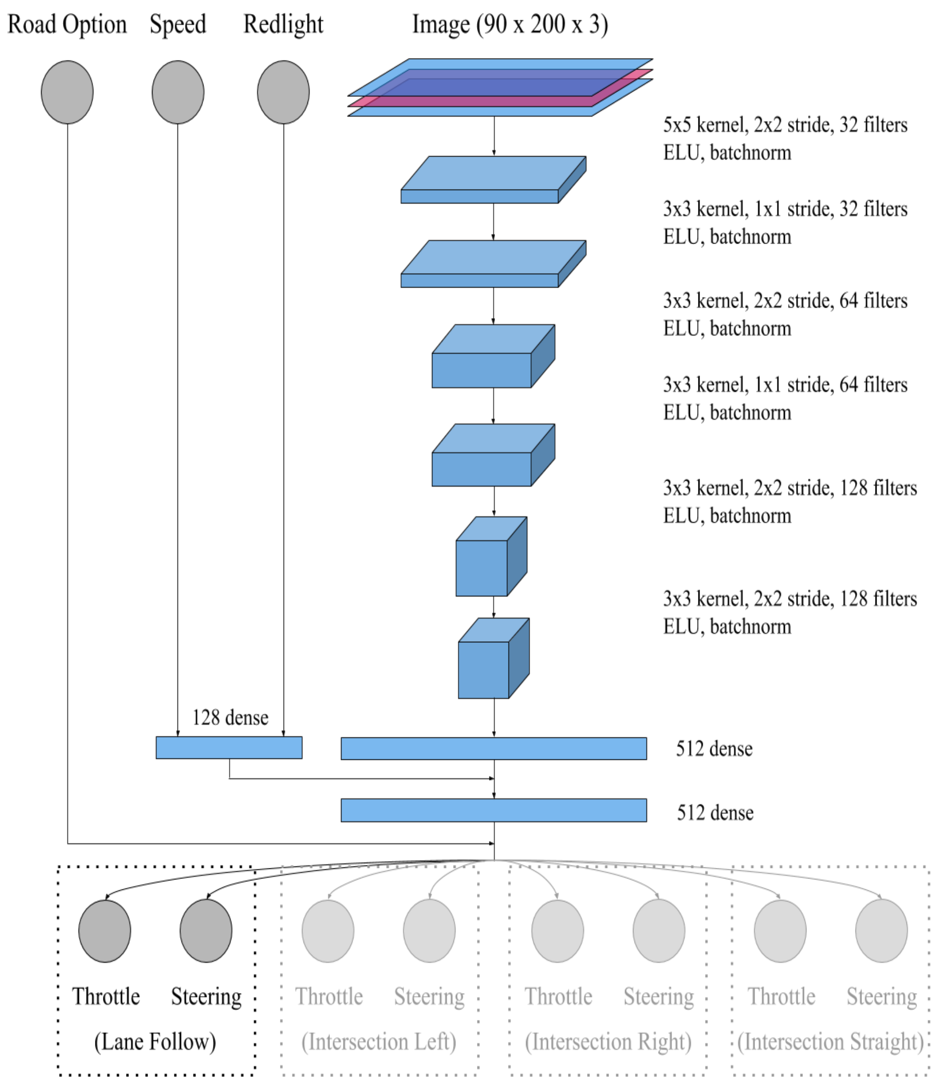

- Robert Patton, Shang Gao, Spencer Paulissen, Nicholas Haas, Brian Jewell, Xiangyu Zhang, Peter Graf, "Heterogeneous Machine Learning on High Performance Computing for End to End Driving of Autonomous Vehicles", 2020 Society of Automotive Engineers World Congress Experience, April 2020.