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

2021 DOE Vehicle Technologies Office

Annual Merit Review Presentation

June 22, 2021

Overview

Timeline

- Project start: October 1, 2018
- Project end: September 30, 2021
- Percent complete: 75%

Budget

- Total project funding
  - DOE share: 100%
  - Contractor share: 0%
- Funding received in FY 2020
  - $ 2.2M
- Funding received in FY 2021
  - $ 1.8M

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	FY21 Q1	FY21 Q2	FY21 Q3	FY21 Q4
1. Utilize an HPC system for ML and scenario evaluation of CAV driving				
2. Develop algorithms for scalable imitation learning, reinforcement learning, and scenario-based evaluation for CAVs control.				
3. Develop energy efficiency metrics for autonomous vehicles; Demonstrate multi-vehicle cooperative driving; Develop and publish open-source data set for machine training.				
4. Develop demos that showcase the range of research performed to date that leverage machine learning in the perception and control for connected autonomous vehicles given the constraints / objectives of safety, destination arrival, and energy efficiency.				

Approach

Connected Autonomous Vehicles

Current: Mimic Human Behavior

Sensing	Perceiving	Controlling	Communicating
ML to read sensor data	ML to define sensor data into a model of physical world	ML to send control signals	ML to send communication signals

Imitation Learning using data (ORNL) Begin FY19

Reinforcement Learning using simulation (NREL) Begin FY19

Imitation & Reinforcement Learning using simulation (ORNL / NREL) Begin late FY19

Scalable HPC-based ML & Simulation Framework for CAV Research (ORNL / NREL) Begin FY20 & Beyond

Future: Machine Behavior

ML spans the spectrum

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

- Synthetic Image Data Generation
- AI trained with Synthetic Data
- Quantitative Evaluation for Driving Metric
- Adversarial Testing & Evaluation
- Running Reinforcement Learning (RL) algorithms at scale (10s of nodes) using high-fidelity CARLA simulator
- RL with in-house medium fidelity KRoad simulator – learning non-perception driving steps that require physics but not graphics
- Studying RL for merging and parking assignment– cooperative and infrastructure guided driving
- Studying feasibility of fusing neural networks using initial common interfacing states

Synthetic Image Data Generation

- Real world data collection: 1) is resource intensive, 2) doesn't provide great coverage of unusual corner cases, and 3) doesn't have labeled truth data unless by very onerous manual means

- Initial version of data set:
  - ~23M images / 507 GB includes RGB, Semantic Segmentation, & LIDAR
  - 3 towns in CARLA (Town 1, 3, and 4)
  - 8 combinations of weather and sun position
  - Weather: Clear, Hard Rain, Soft Rain, and Wet
  - Sun position: Noon, Sunset
  - 2 hours of driving per combination per town
  - No pedestrians / No traffic; focused on static objects

- Next version
  - Include dynamic objects (e.g., pedestrians, vehicles)
  - Include DVS camera & bounding boxes
  - Release via Livewire

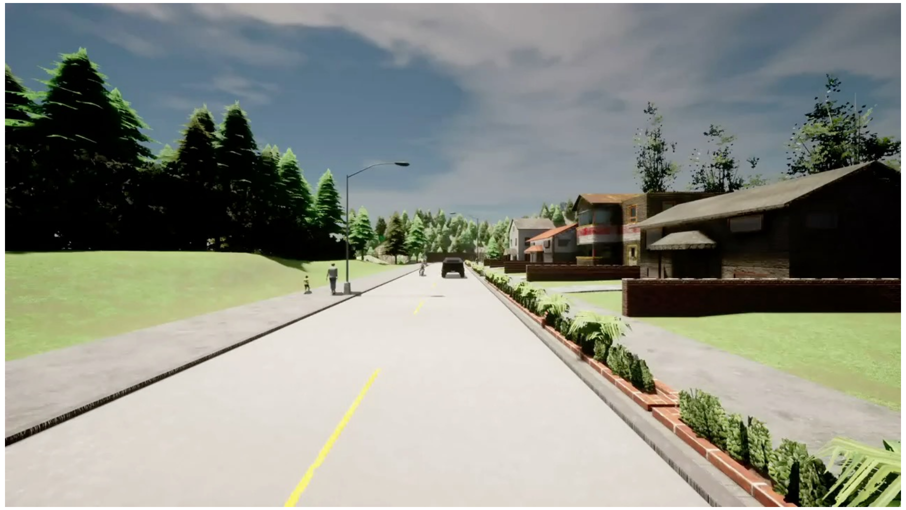

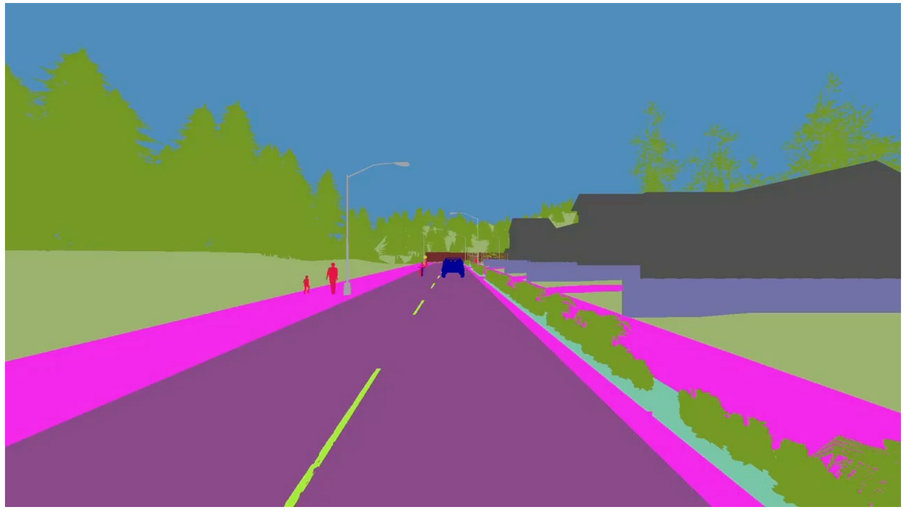

OAK RIDGE
National Laboratory

AI trained with Synthetic Image Data

- 622K images subset from an 8-hour data set
- 28 training epochs
- ~16 hours to train

- 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

Input

Output

Brake
Steering
Throttle

SSNet

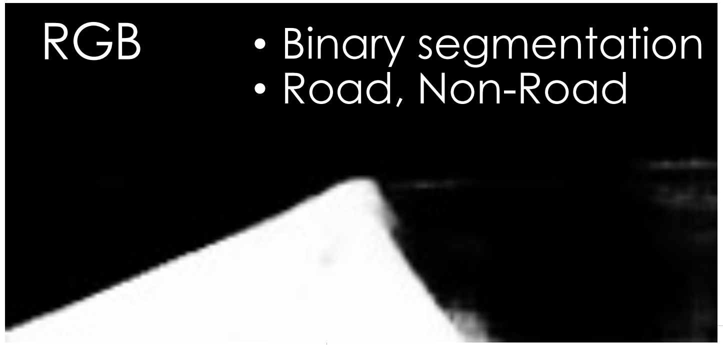

- Binary segmentation
- Road, Non-Road

SSNet: Multi-task network that maps RGB images & LIDAR data to a set of brake, steering, and throttle values as well as a semantic segmented image

QED - Quantitative Evaluation for Driving Metric

- Tiered approach allows for focused evaluation of specific behaviors
- ORNL has developed a Tier 1 metric that measures the precise path that the car follows and gives a score from 1-100: 1 being the worst, 100 being the best
- Captures various types of driving behavior including distance driven, staying in lane, weaving, speeding, and collisions
- We compared QED to the scores that humans give:
  - QED provides a ranking of scores that are consistent and indistinguishable from human rankings
- *A quantitative metric now makes adversarial testing possible*

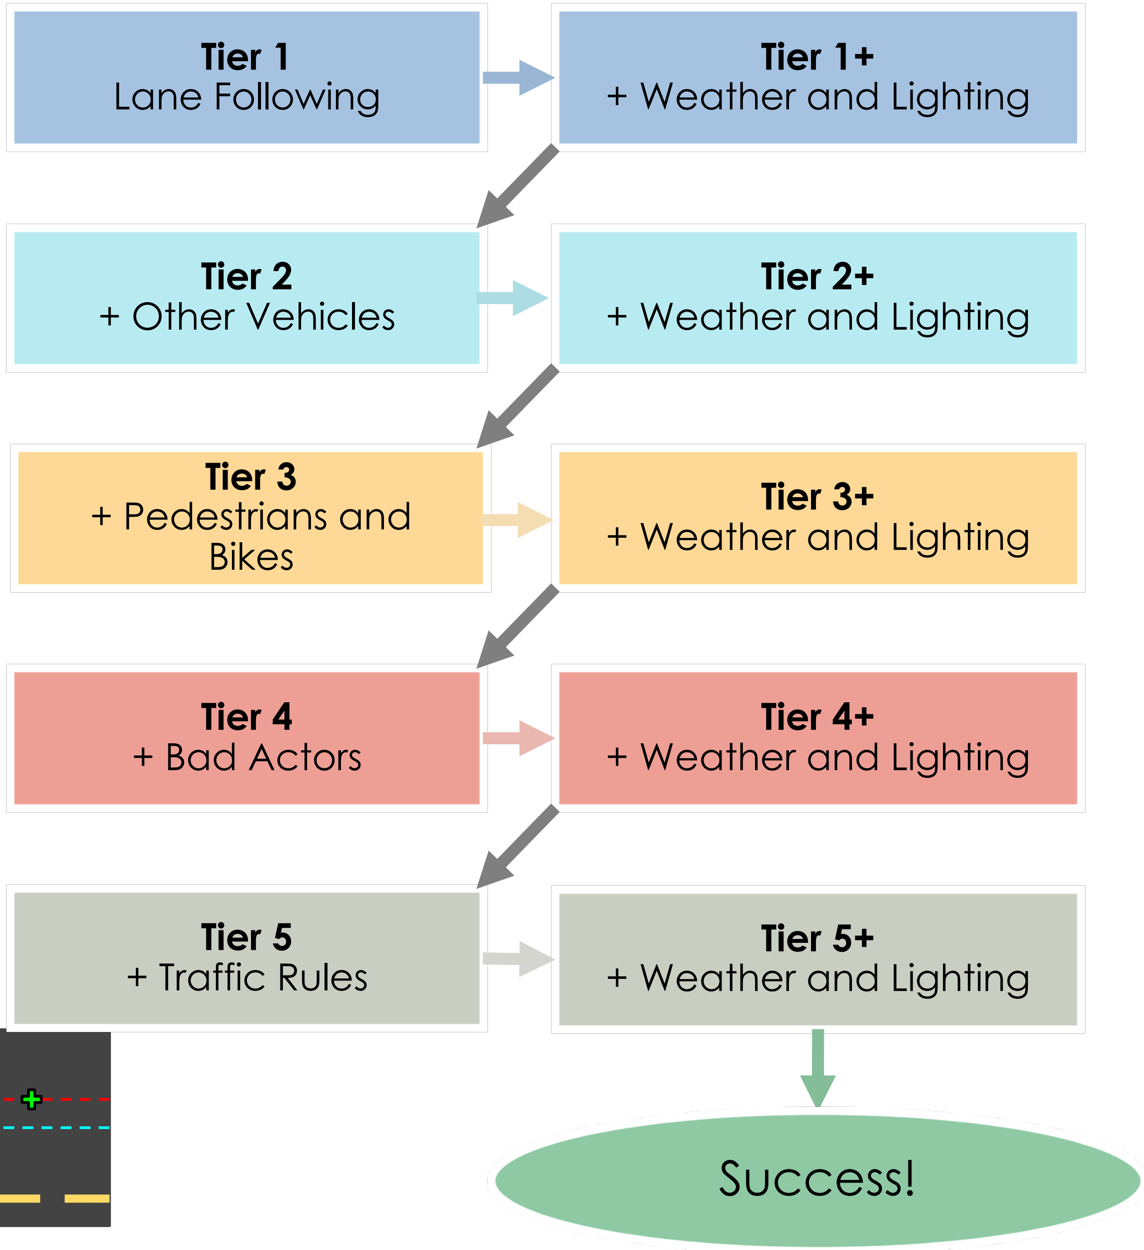

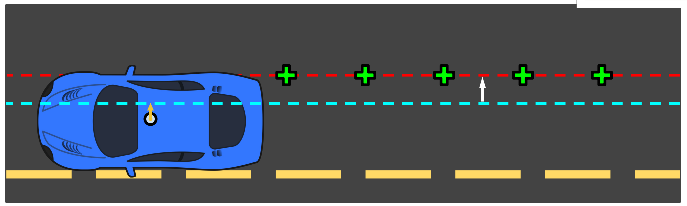

OAK RIDGE National Laboratory

Adversarial Testing & Evaluation

- Gremlin finds the scenarios that cause the autonomous vehicle to fail
- 2020 – 2021 DOE ALCC Proposal: Evolutionary Multi-scenario Simulation Environment for Autonomous Vehicle Testing
  - Partnered with GM; 250K compute hours on ORNL's Summit
- Running on Summit!
  - Several 100 node runs completed; larger runs planned
  - Initial results from a 100 node runs shown on the right
  - Red dots on road map shows where Gremlin finds a problem

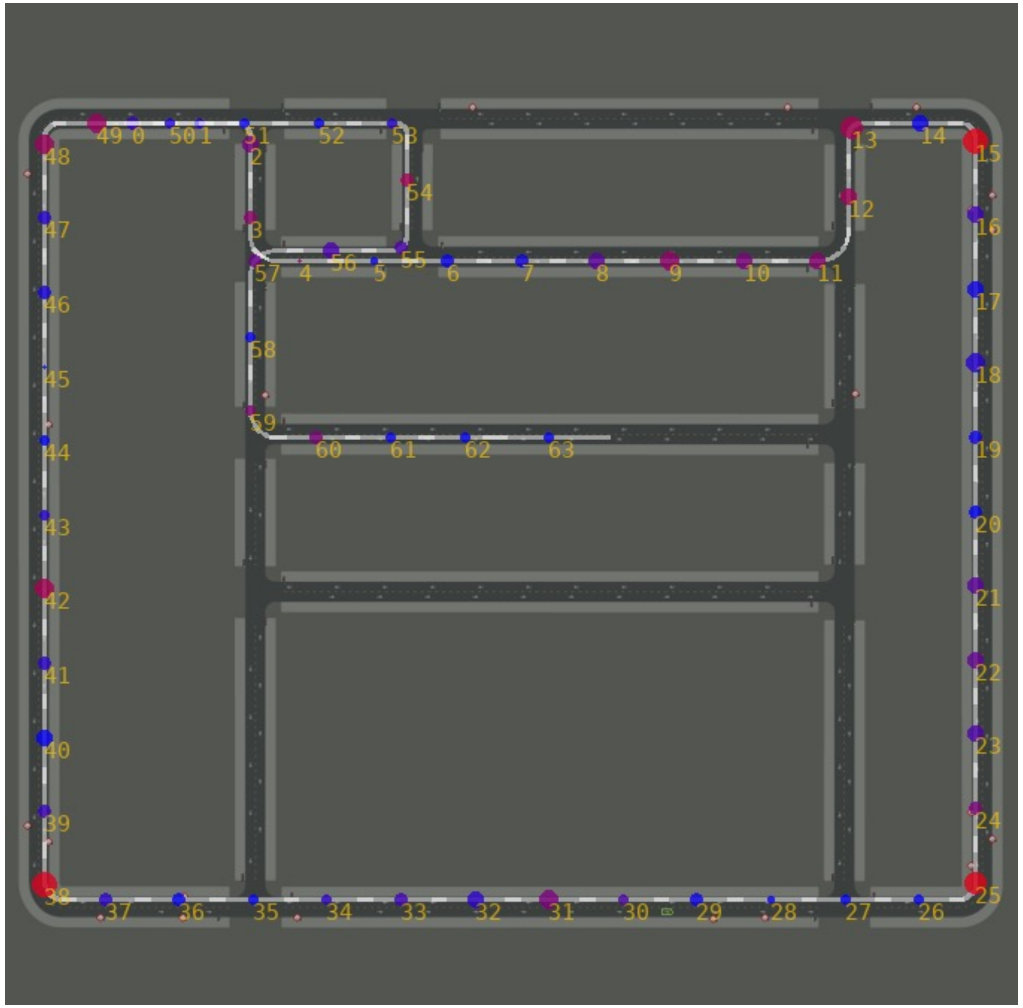

Reinforcement Learning (RL) with CARLA – scaling RL

Scaling CARLA learning on HPC

- **[Scaling vehicles]** Reducing simulation overhead by leveraging CARLA's client-server architecture: parallel driving experience collection via multiple clients/vehicles within a single CARLA world/server.
- **[Scaling simulations]** Fully utilizing the computing power of a single HPC node: running four* CARLA servers in parallel per node. (*Determined by GPU memory size).
- **[Scaling nodes]** Multiple HPC nodes are used to further scale up the learning. Multi-node learning is enabled by using a distributed computing platform called Ray and a RL library based on it called RLlib.

See illustration to the right for the overall structure.

Why is scaling RL so hard?

- RL dynamically gathers new data
- Uses it to learn a single policy
- Frequent communication required
- Leads to bottlenecks to scaling

Solution will be hierarchical (e.g., multiple learners) and distributed. This is a subject of ongoing research.

CARLA HPC Multi-node Learning Architecture

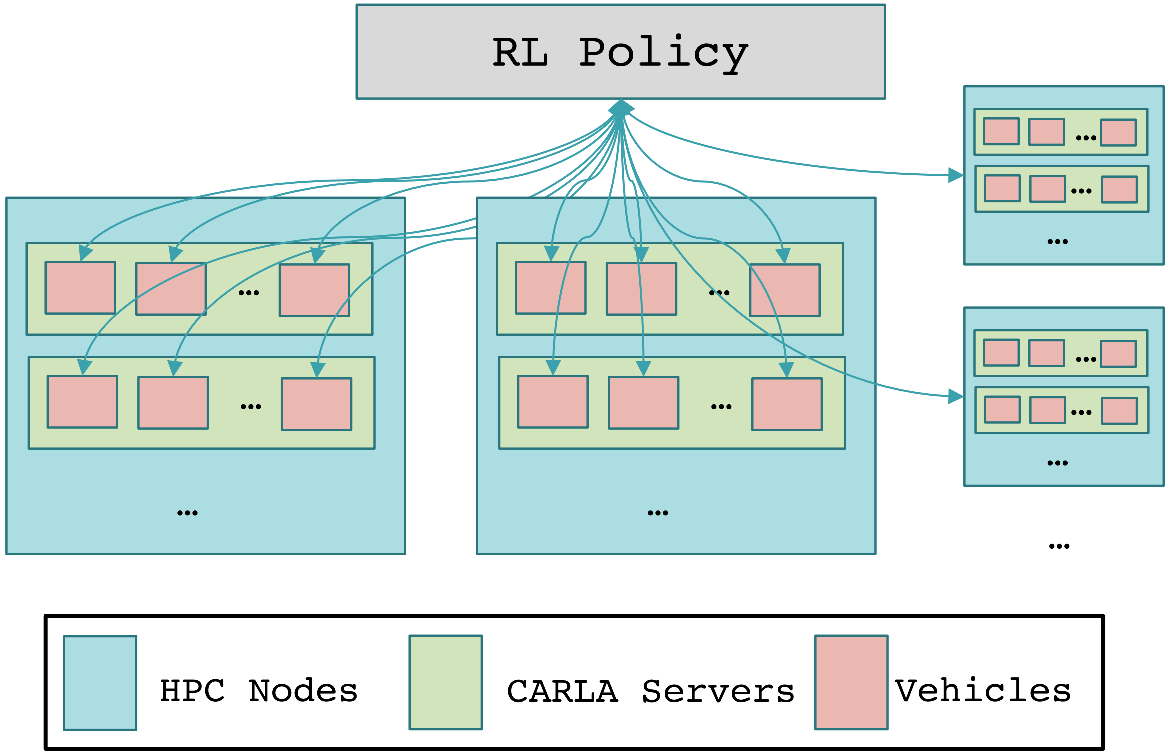

To date:
20 HPC Nodes
x 4 CARLA Simulation Servers Per Node
x 16 Vehicles Per Server = 1280 vehicles

RL for path tracking

Many sub-steps in CAVs control pipeline require physics but not perception. NREL developed KRoad to fill this need.

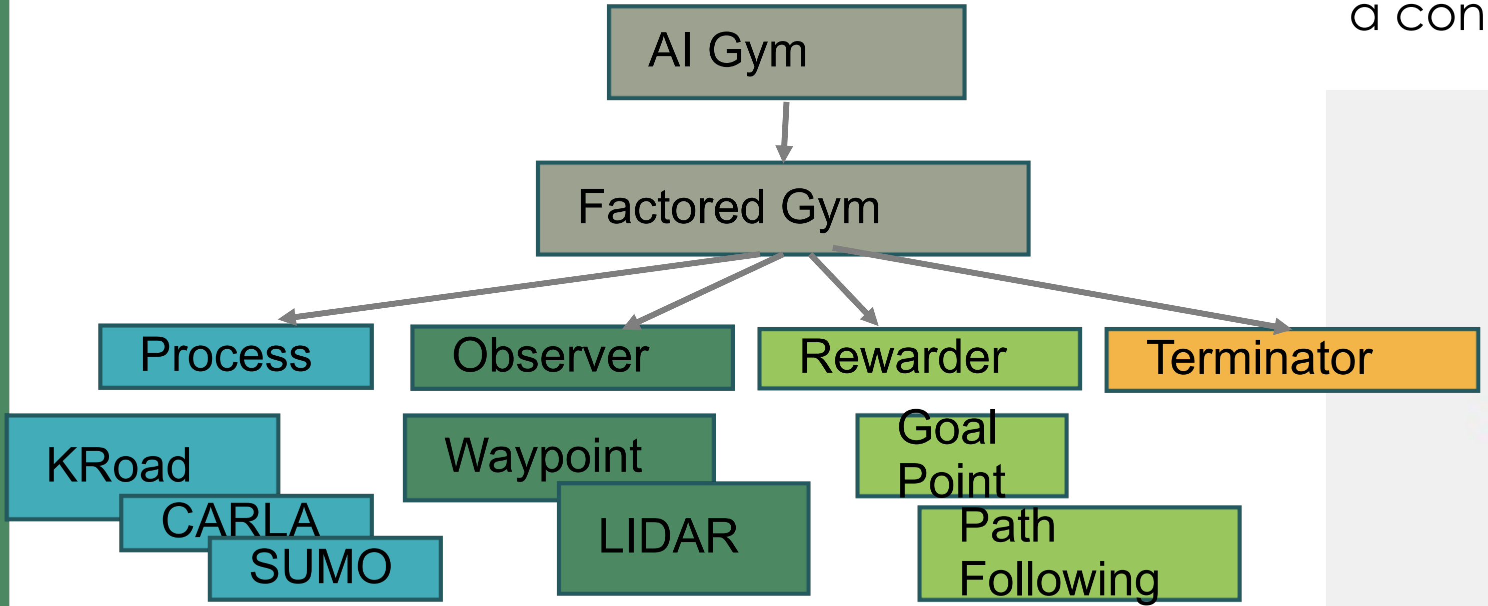

KRoad is a highly modular AI gym that implements the physics of the dynamic bicycle model and allows for a wide variety of physics-based reinforcement learning research for vehicle controls.

Actuating a vehicle to follow a given path is still a challenging problem. We used KRoad to train a controller with RL for path tracking

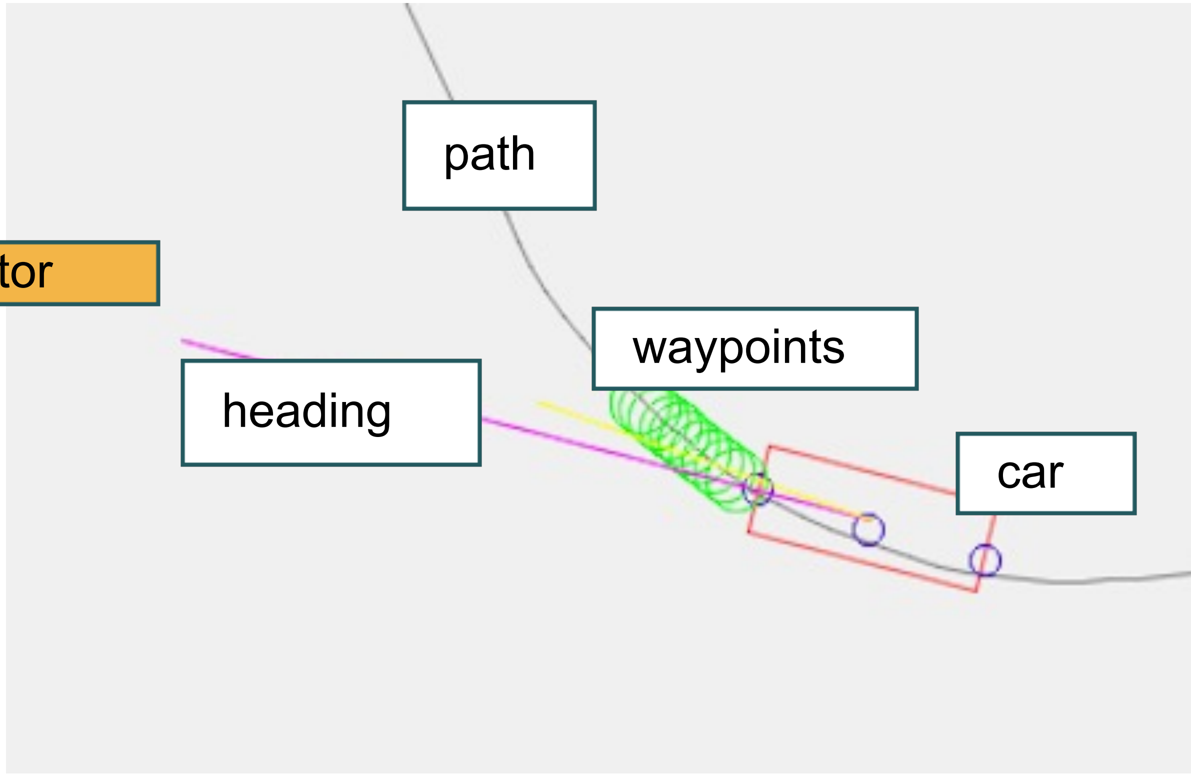

Trained RL agent navigates random path

RL for cooperative and infrastructure guided driving

Cooperative lane changing

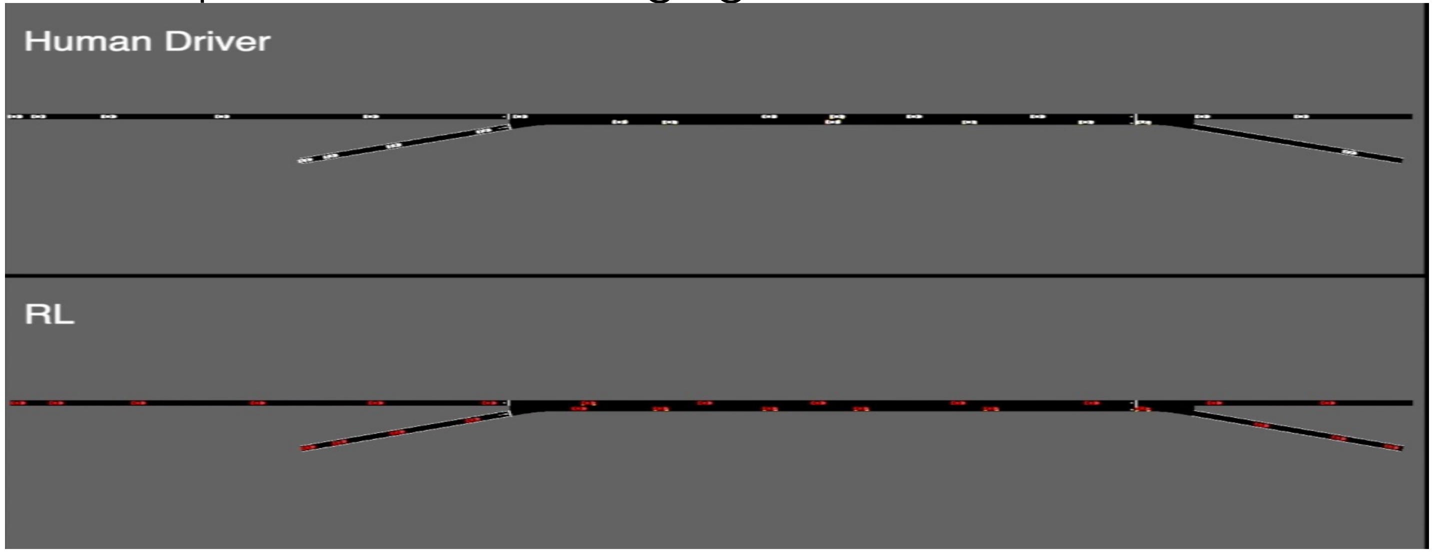

	MPG	CO2 (g/mi)	NOx (mg/mi)	Stops	Speed (m/s)	Outflow (veh/h)
Human Driver	13.0	669.8	282.3	14.2	6.7	1996
RL	14.9	586.6	245.8	2.7	14.5	2407
Difference	+15%	-12%	-13%	-81%	+116%	+21%

RL controller improves on baseline on several important measures

Infrastructure-based speed control

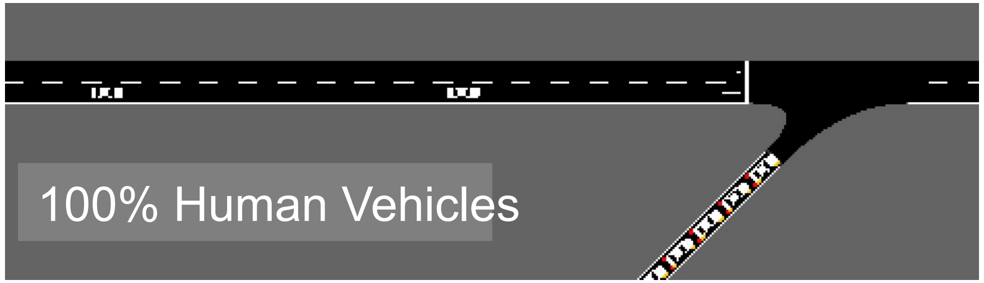

Without RL, merge lane bottleneck

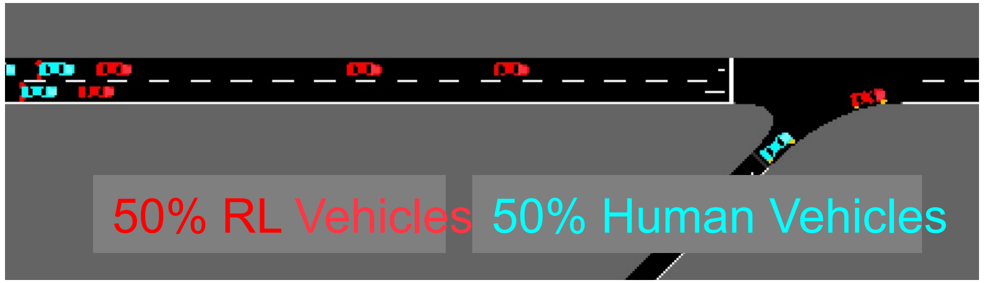

RL learn speeds that allow space to merge

Cooperative curbside parking

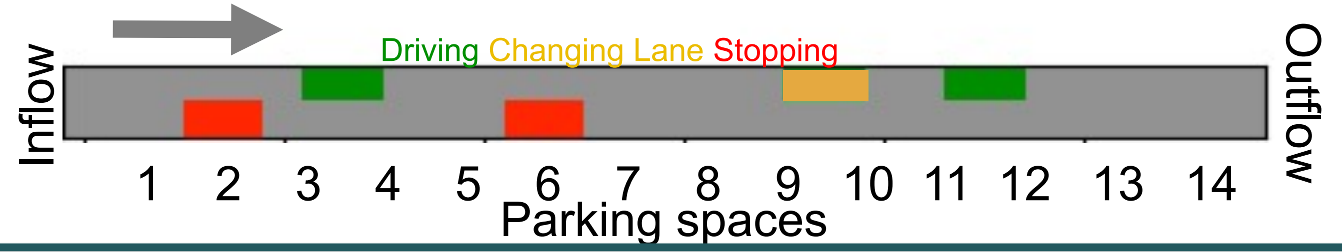

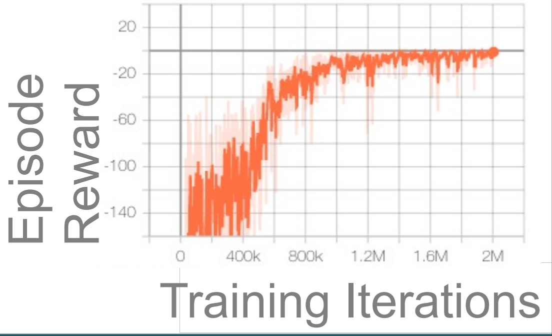

NREL
NATIONAL RENEWABLE ENERGY LABORATORY

Fusing Neural Networks for Path Planning and Path Following

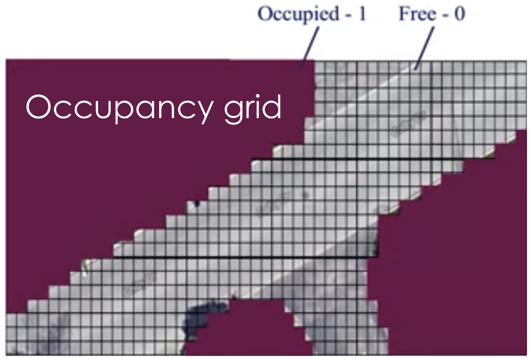

The occupancy grid provides a map of free vs occupied space in a grid around the ego vehicle.

Path Planning

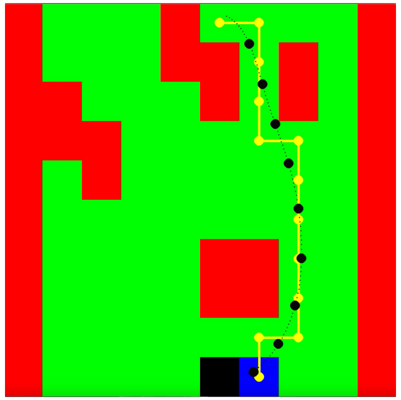

We have solved path planning using the occupancy grid as input in a novel way, using RL to solve the “maze-like” part of the problem with crude left/right/up/down steps, then smoothing the path with spline interpolation. The output is the set of black waypoint dots from top to black target square at bottom.

Waypoint Following

Our RL solution for waypoint following is the subject of previous work; the vehicle follows the waypoints produced by the path planning network

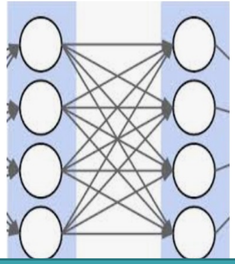

Path planning

+

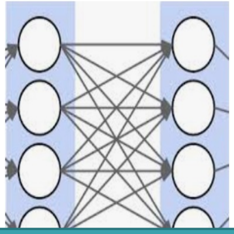

Waypoint Following

=

Full solution

Research questions:

Is learning each step separately more efficient than trying to learn them all at once?

When 2 networks are fused, what happens to “latent” space in the resulting middle?

Responses to Previous Years Reviewers Comments

- "So far, it is not very clear how to quantify a driver to be an expert for imitation learning. Also, it is not very clear if there would be any other side effects for the entire traffic flow (under different CAV penetration rates, including the extreme case of 100% CAVs) if autonomous driving is trained in this way."
  - *Response*: A quantitative evaluation metric was developed (slide 9) to quantify the driving behavior of the machine. This same metric has been used to assess driving behavior of a human driver within CARLA. In March 2020, CARLA released in version 0.9.8 and experimental connection with SUMO. We have begun investigating this connection and how changes at the individual vehicle level within CARLA could be further studied at the traffic level with SUMO. We are using CARLA version 0.9.8 on ORNL's Summit supercomputer for this work.
- "The reviewer thought this will take time to figure out because the replication of the human driver in simulation is needed to then assess the energy savings."
  - *Response*: A human driver could use CARLA to drive virtually with all of the data from that simulation stored and used for analysis and comparison to the machine. Furthermore, there are models for human drivers within SUMO that could be used as well if the energy savings assessment is done at the traffic level.

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?
  - Overcome inherent bottlenecks of scaling RL
- **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)?
- **Energy Efficient AI:**
  - How can AI be developed to not require power hungry computing?
- **Connectivity:**
  - How do changes to driving behavior at the vehicle (micro) level affect changes to traffic flow (macro level)?
  - How does connectivity to vehicles / infrastructure at the vehicle (micro) level affect change at the traffic (macro) level?

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):**
  - Synthetic Image Data Generation
  - AI trained with Synthetic Data
  - Quantitative Evaluation for Driving Metric
  - Adversarial Testing & Evaluation
  - Running Reinforcement Learning (RL) algorithms at scale using high-fidelity CARLA simulator
  - RL with in-house medium fidelity KRoad simulator – learning non-perception driving steps that require physics but not graphics
  - Studying RL for merging and parking assignment– cooperative and infrastructure guided driving
  - Studying feasibility of fusing neural networks using initial common interfacing states
- **Future Work:**
  - Energy Efficient AI
  - Scaling and integration of tools supporting Reinforcement Learning for driving simulators
  - Scaling up the computing resources and scenario generation

QUESTIONS?

Technical Backup