Learning a Robust Multiagent Driving Policy for Traffic Congestion Reduction

Yulin Zhang, William Macke, Jiaxun Cui, Daniel Urieli, Peter Stone

Learning Agent Research Group
The University of Texas at Austin
General Motors R&D Labs
Sony AI

Traffic congestion caused by the stop-and-go waves

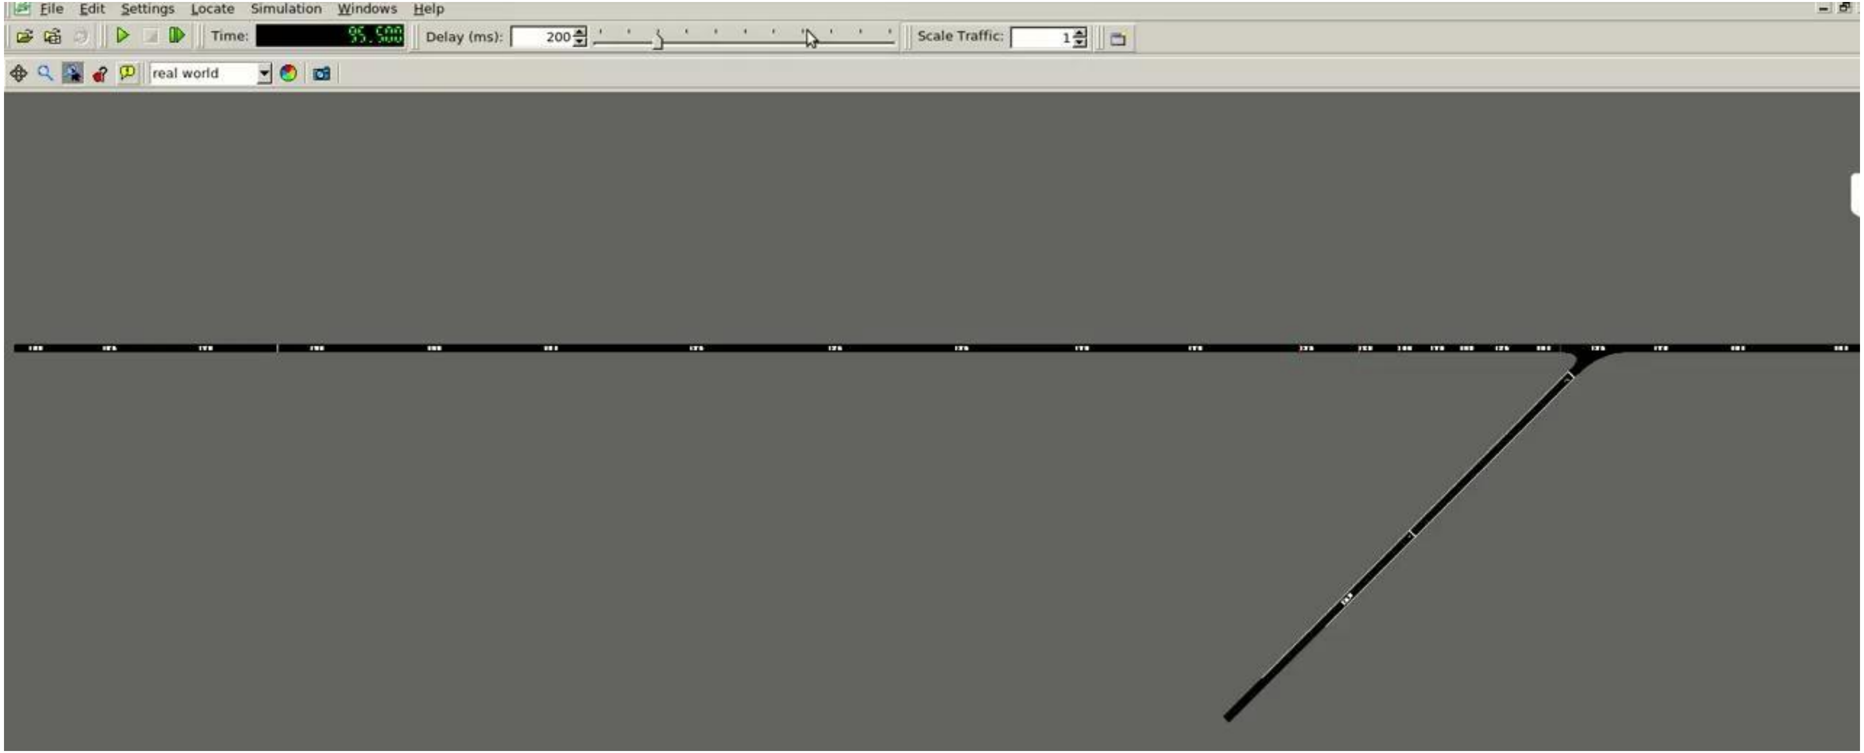

Training Autonomous Vehicles to Reduce Traffic Congestion

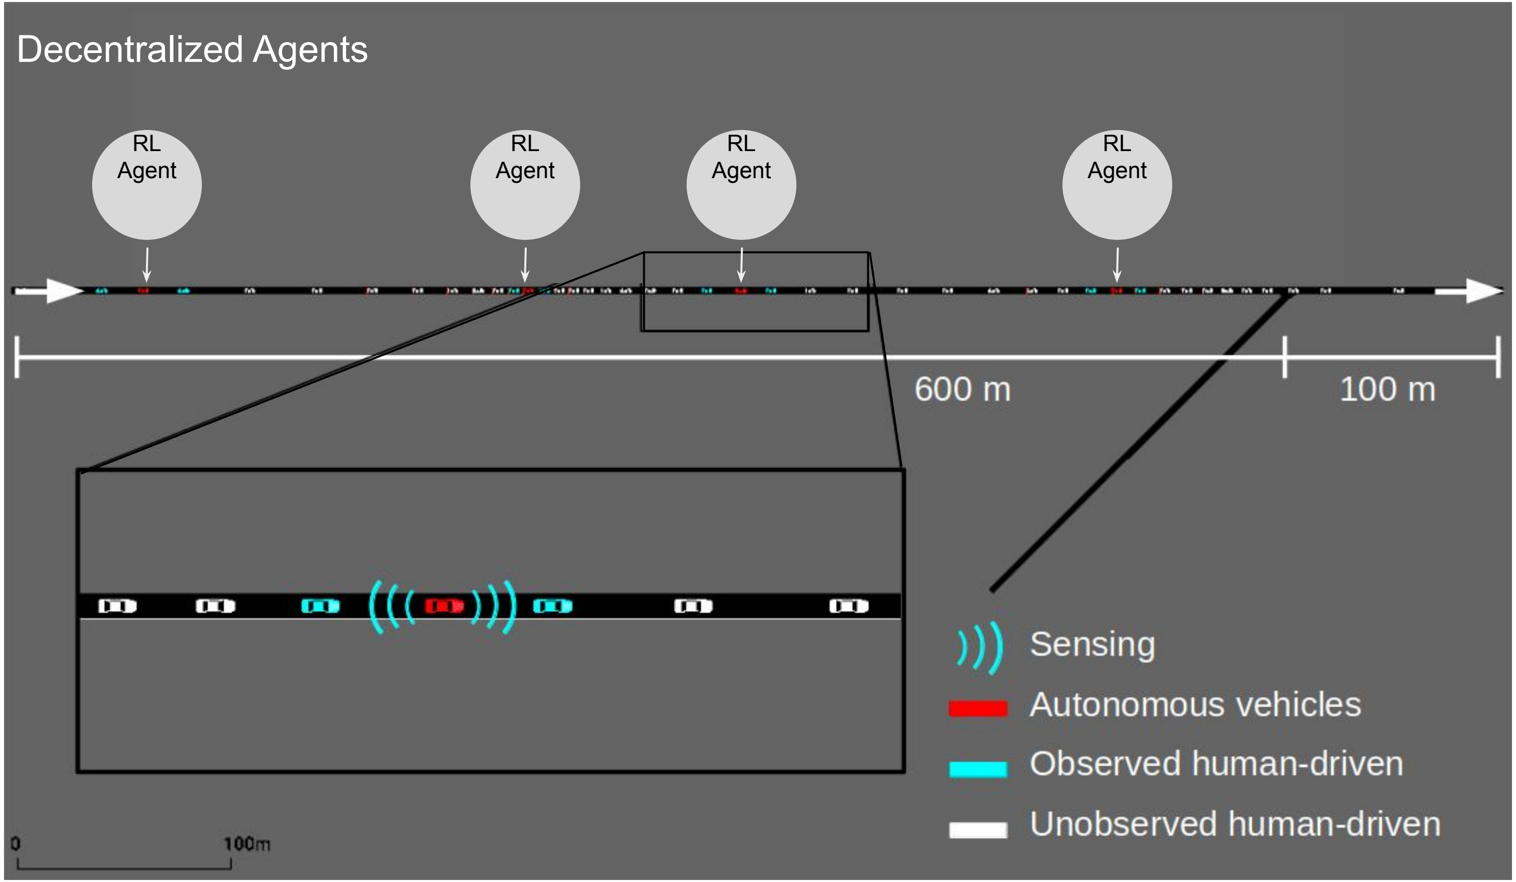

Training Autonomous Vehicles to Reduce Traffic Congestion

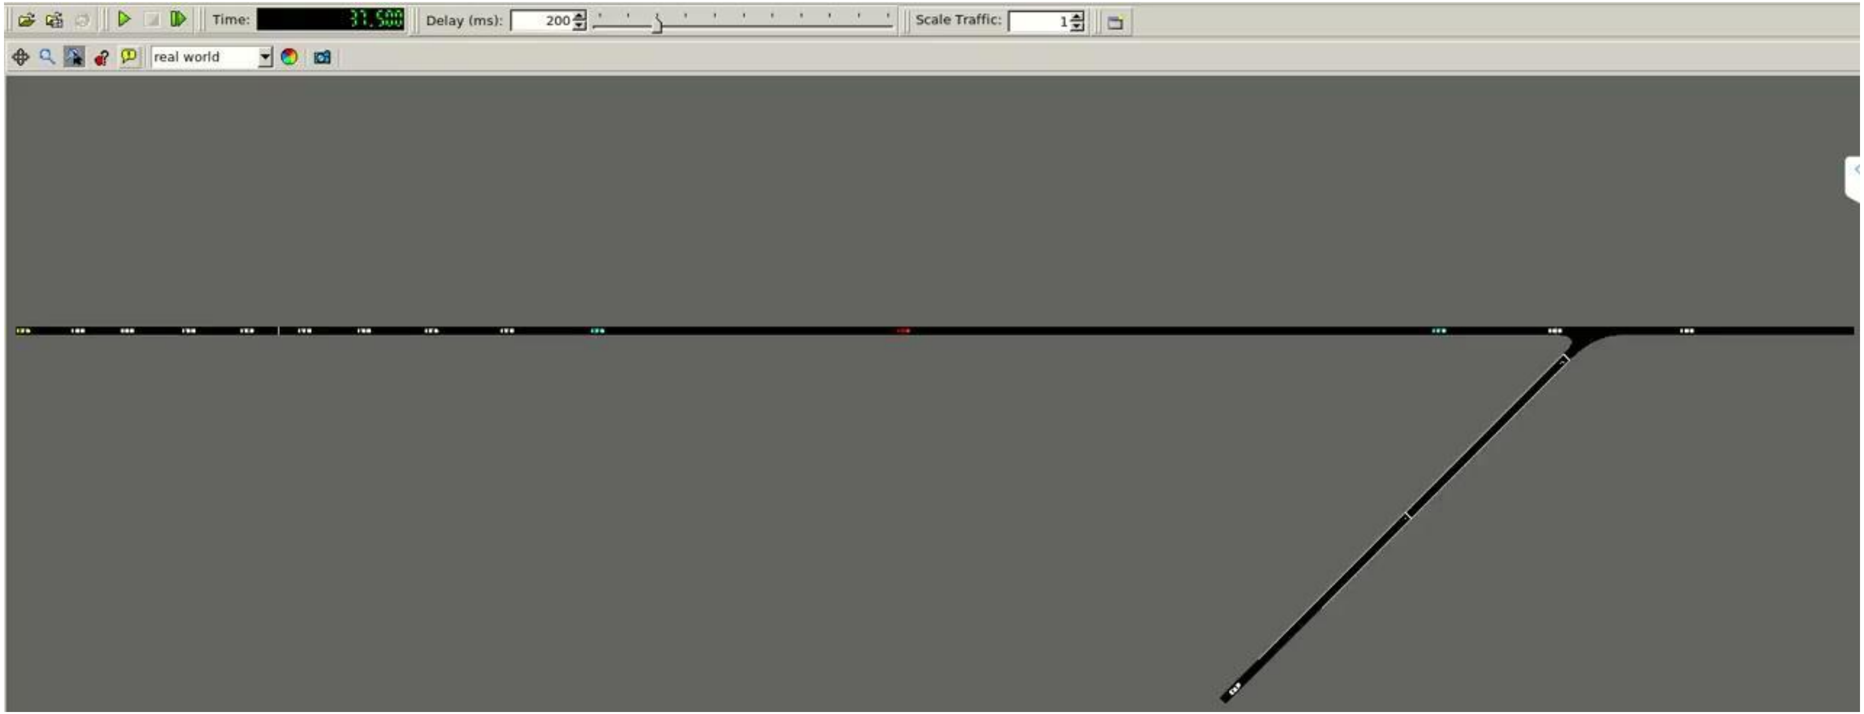

But we are still far from real-world deployment

Policies were trained and tested under similar conditions

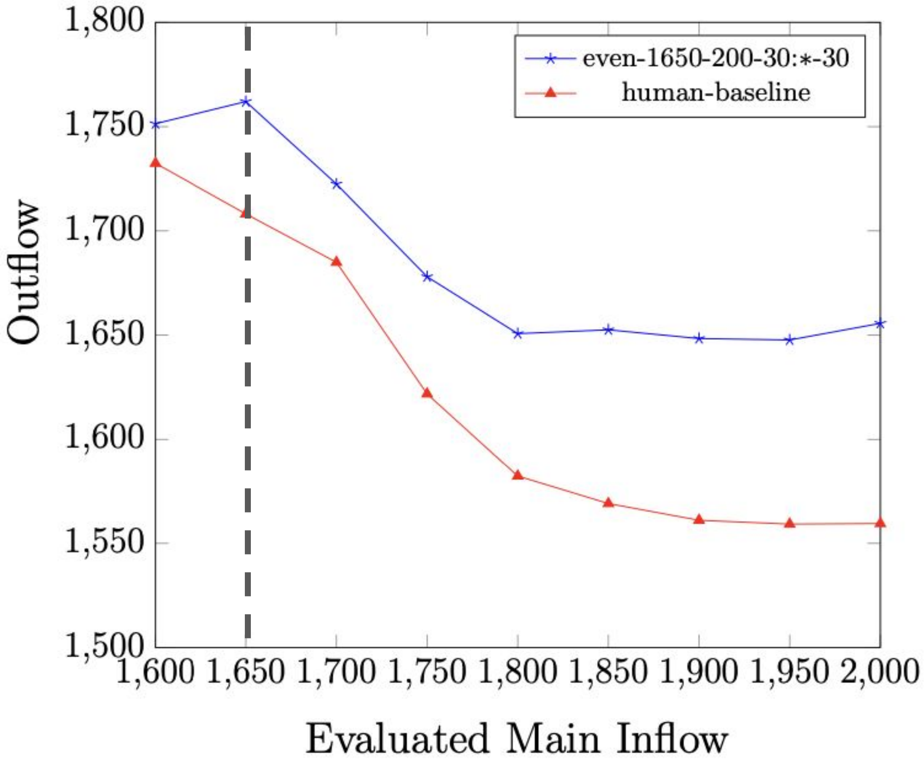

Progress to reduce traffic congestion using AVs

- Traffic reduction in both closed network (circular roads) and open network [Wu et al., 2017b; Kreidieh et al., 2018; Vinitsky et al., 2018].
- Centralized and Decentralized driving policies [Vinitsky et al., 2018; Cui et al., 2021].
- Developing robust driving policies:
  - The robustness of a hand-coded policy is examined over different AV penetration and driving aggressiveness [Parvate, 2020].
  - Generalizing to different traffic densities on a closed ring road [Wu et al., 2021].
  - Negative results on the generalization of a single-lane policy to a double-lane ring road [Cummins et al., 2021].
  - Robust policy is developed in bottleneck scenario [Vinitsky et al., 2020].

In this paper,

- We first develop a single-lane decentralized policy that is robust to:
  - AV placement in traffic
  - Traffic flow
  - Fraction of AVs in traffic (AVP)
- We demonstrate that this is also robust to different road geometry:
  - Road with two merging ramps

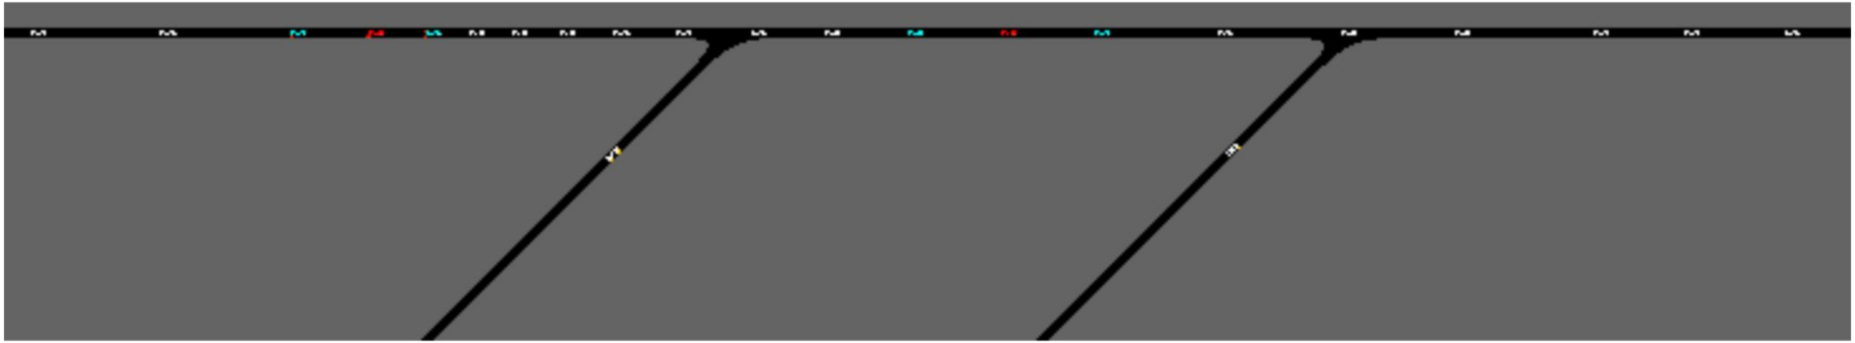

In this paper,

- We first develop a single-lane decentralized policy that is robust to:
  - AV placement in traffic
  - Traffic flow
  - Fraction of AVs in traffic (AVP)
- We demonstrate that this is also robust to different road geometry:
  - Road with two merging ramps
  - Double-lane road

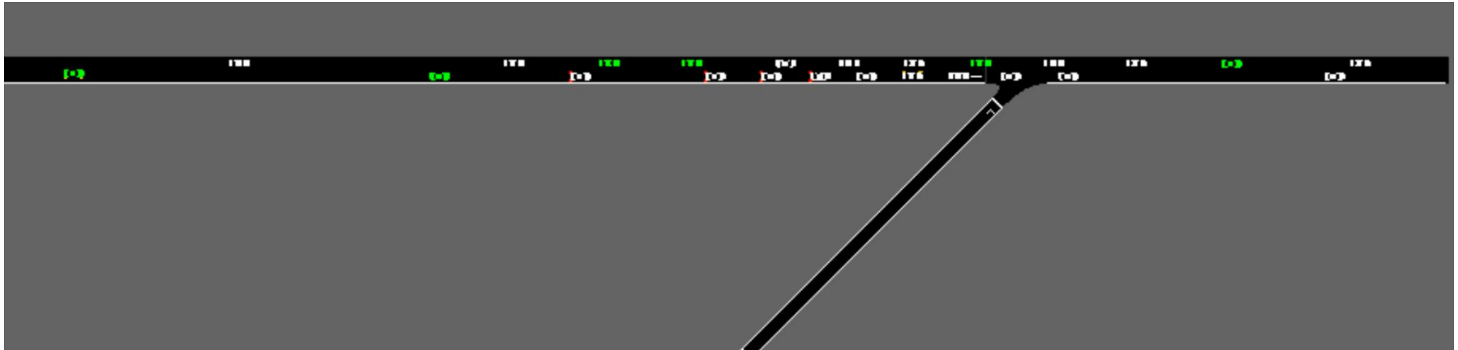

Single-lane decentralized policy: vehicle placement

even placement

random placement

Policy evaluated under random vehicle placement

Training: even or random vehicle placement, main inflow 2000, train and evaluate at the same AVP,
Evaluating: **random** vehicle placement, main inflow= [1200, 2000]

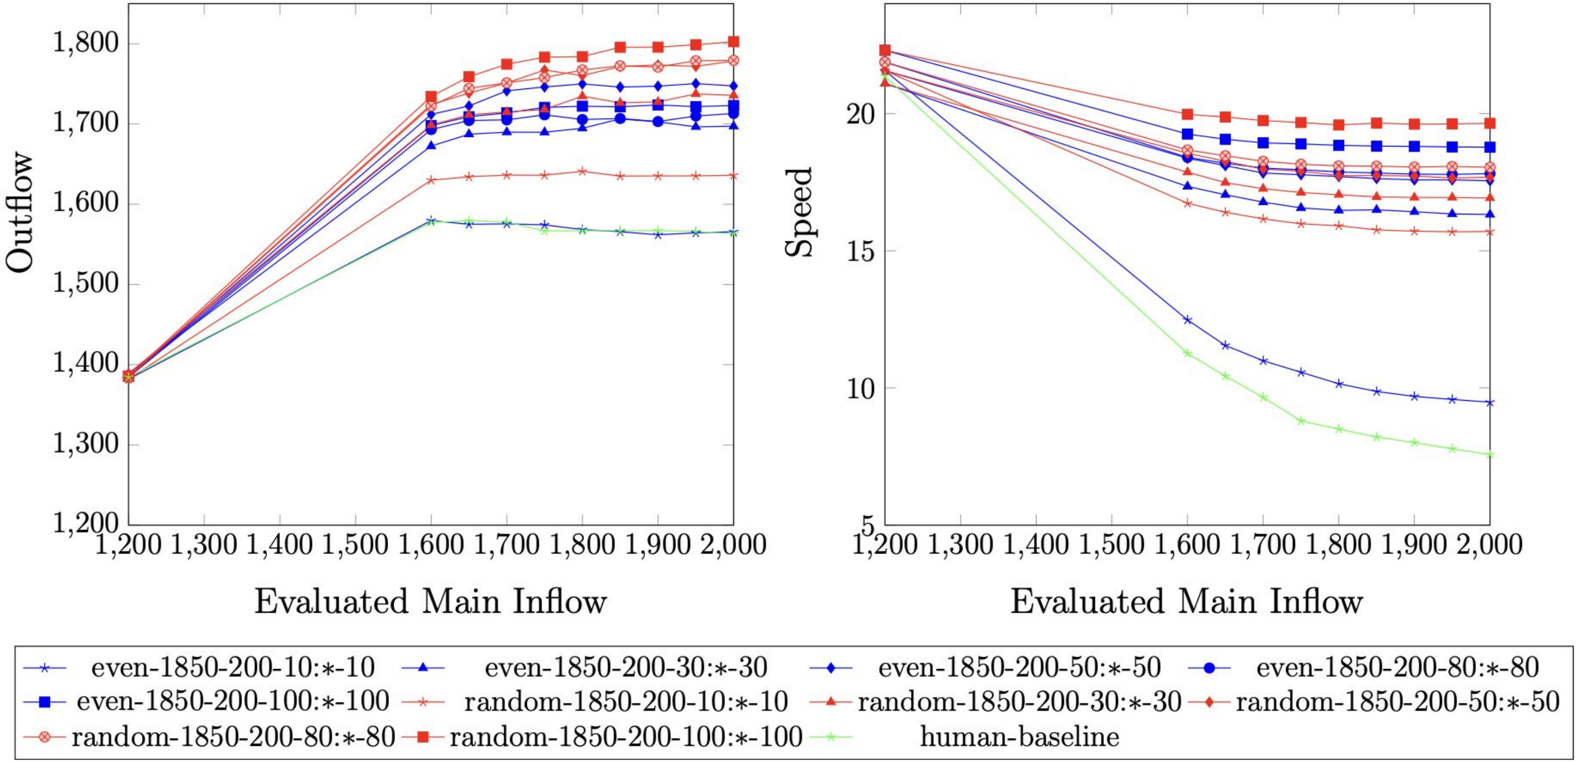

Policy evaluated under even vehicle placement

Training: even or random vehicle placement, main inflow 2000, train and evaluate at the same AVP,
Evaluation: **even** vehicle placement, main inflow= [1200, 2000]

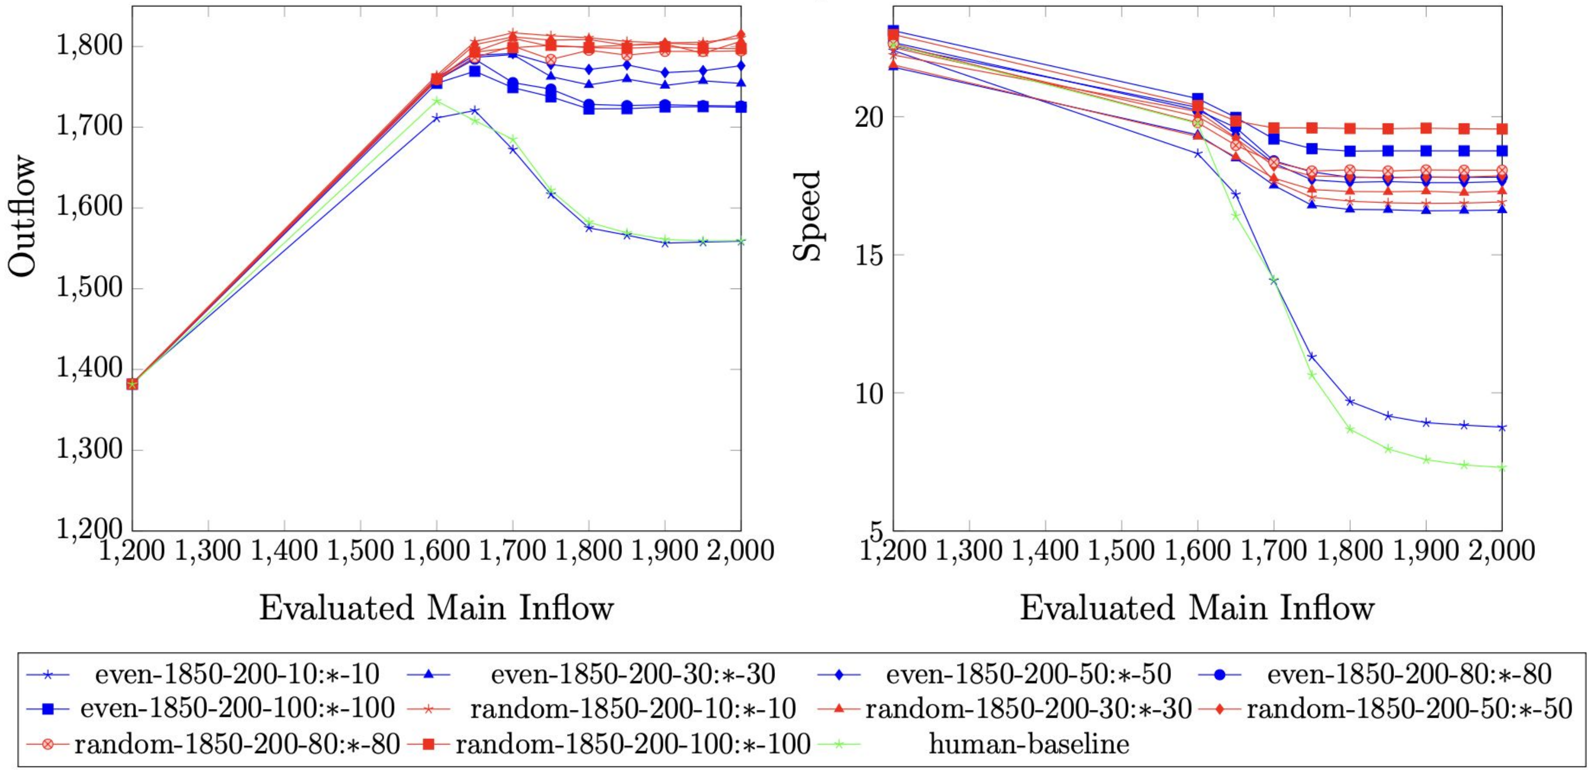

Single-lane decentralized policy: AV penetration/faction

10% AVP

30% AVP

Policy trained under medium AVP (30%)

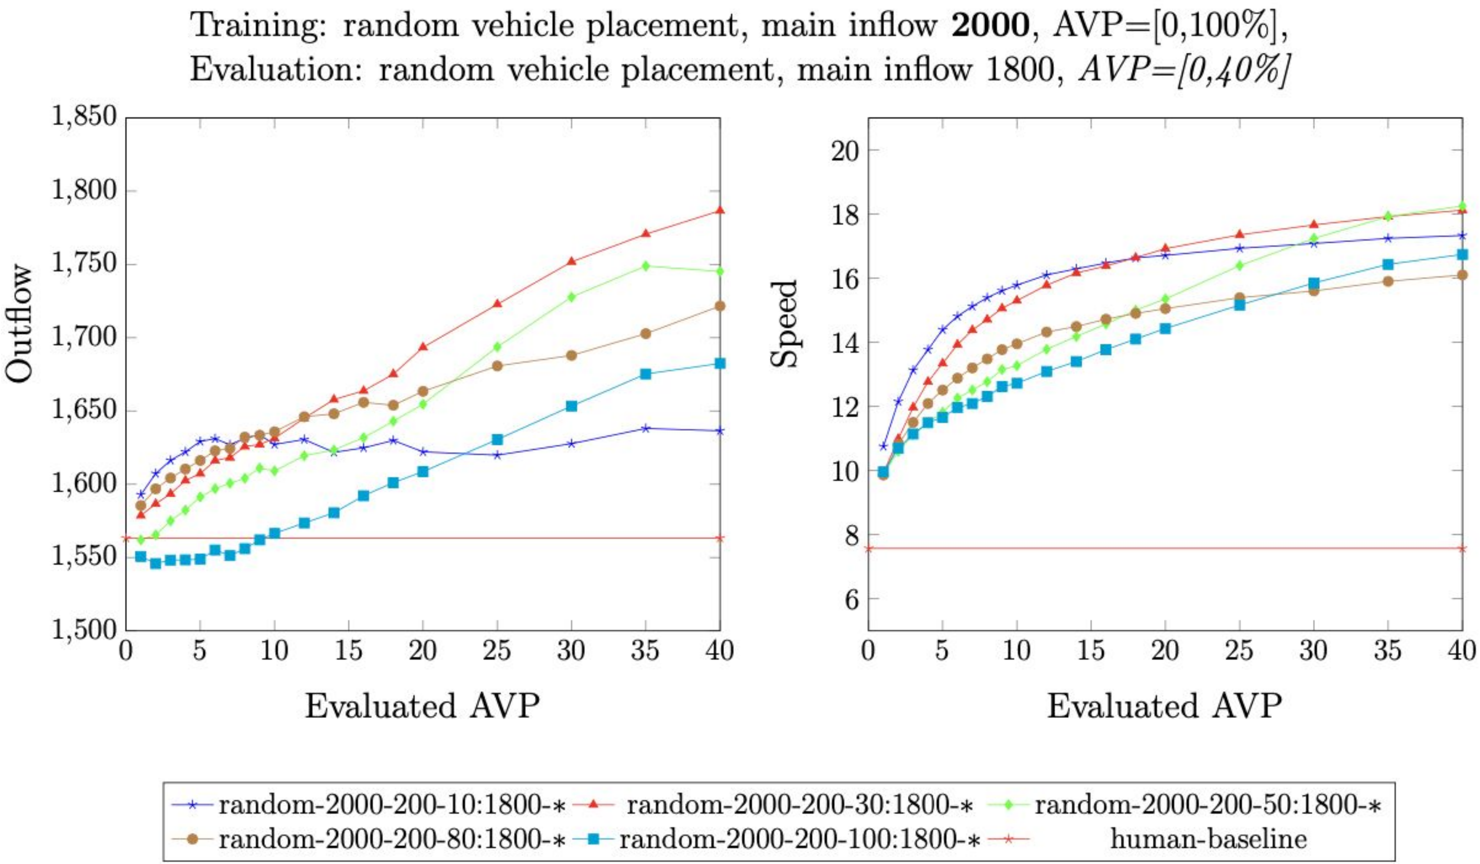

Single-lane decentralized policy: Inflow

1200 veh/hour

2000 veh/hour

Policy trained under high inflow (2000 veh/hour)

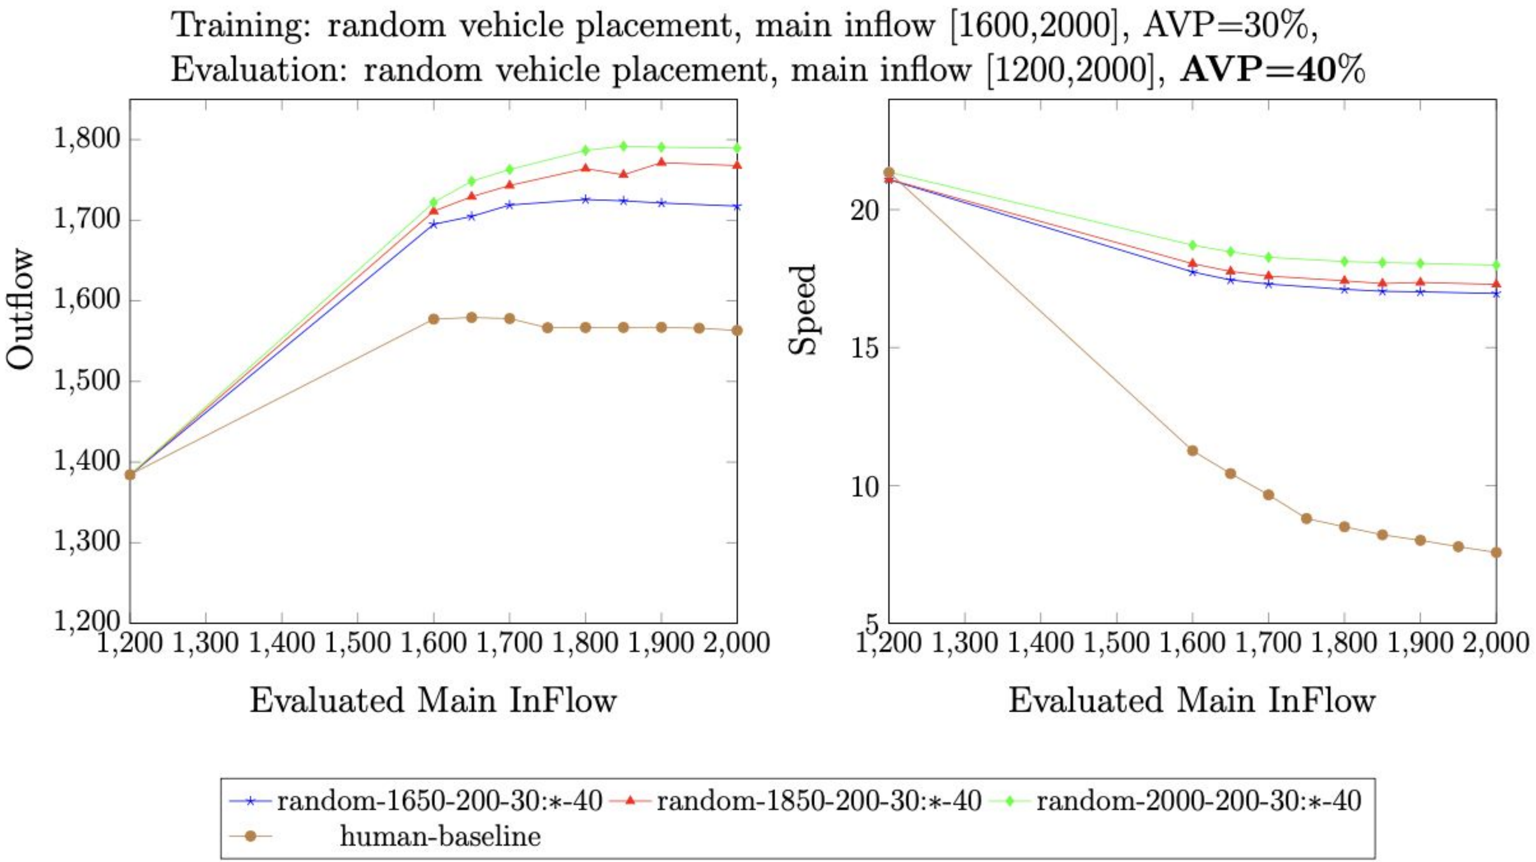

Single-lane decentralized policy: deployed with two ramps

The distance between two ramps is 200 m

The distance between two ramps is 800 m

Single-lane decentralized policy: deployed with two ramps

Training: random vehicle placement, main inflow 1800, merge inflow 200, AVP=30%,
Evaluation: random vehicle placement, main inflow 1800, merge inflow 200 for each ramp, AVP=[0,40%]

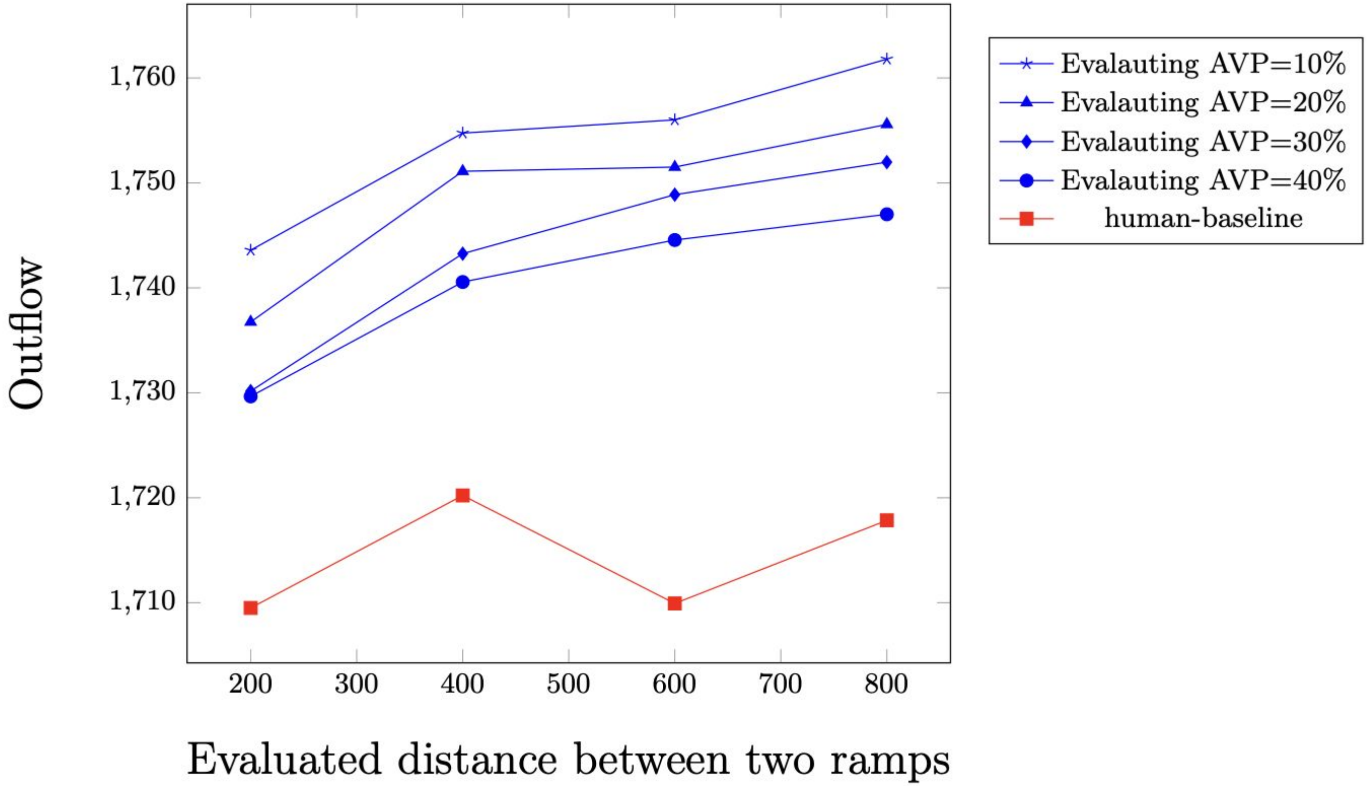

Single-lane decentralized policy: deployed in the right lane

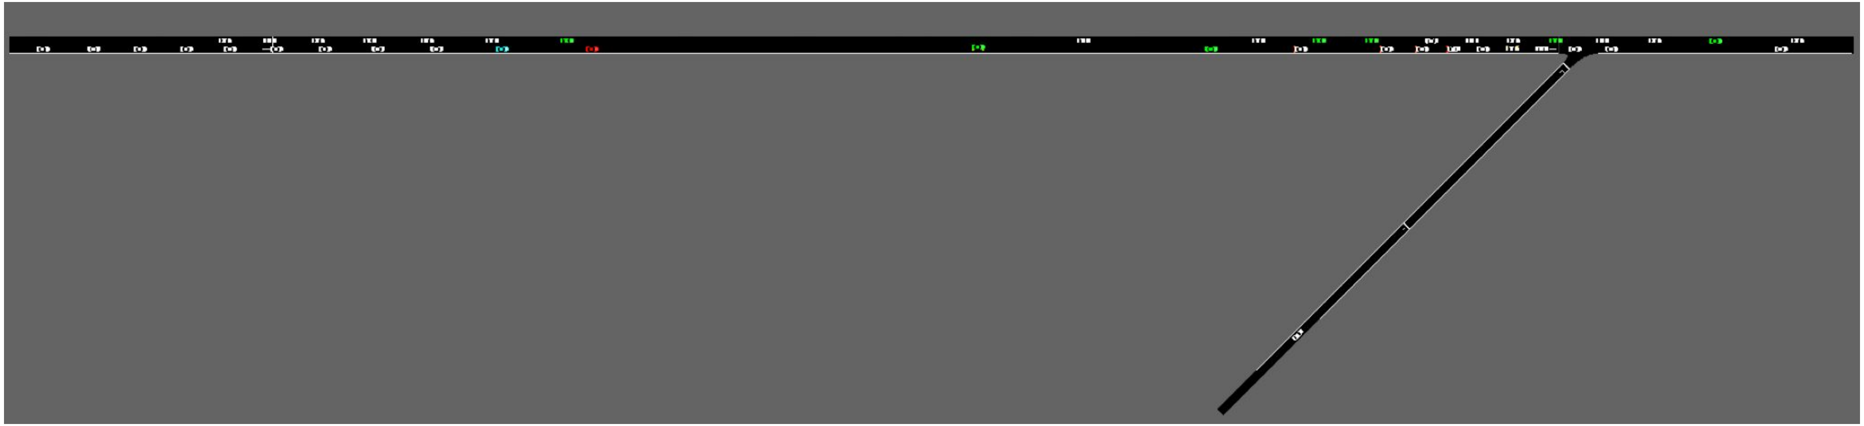

Single-lane decentralized policy: deployed in the right lane

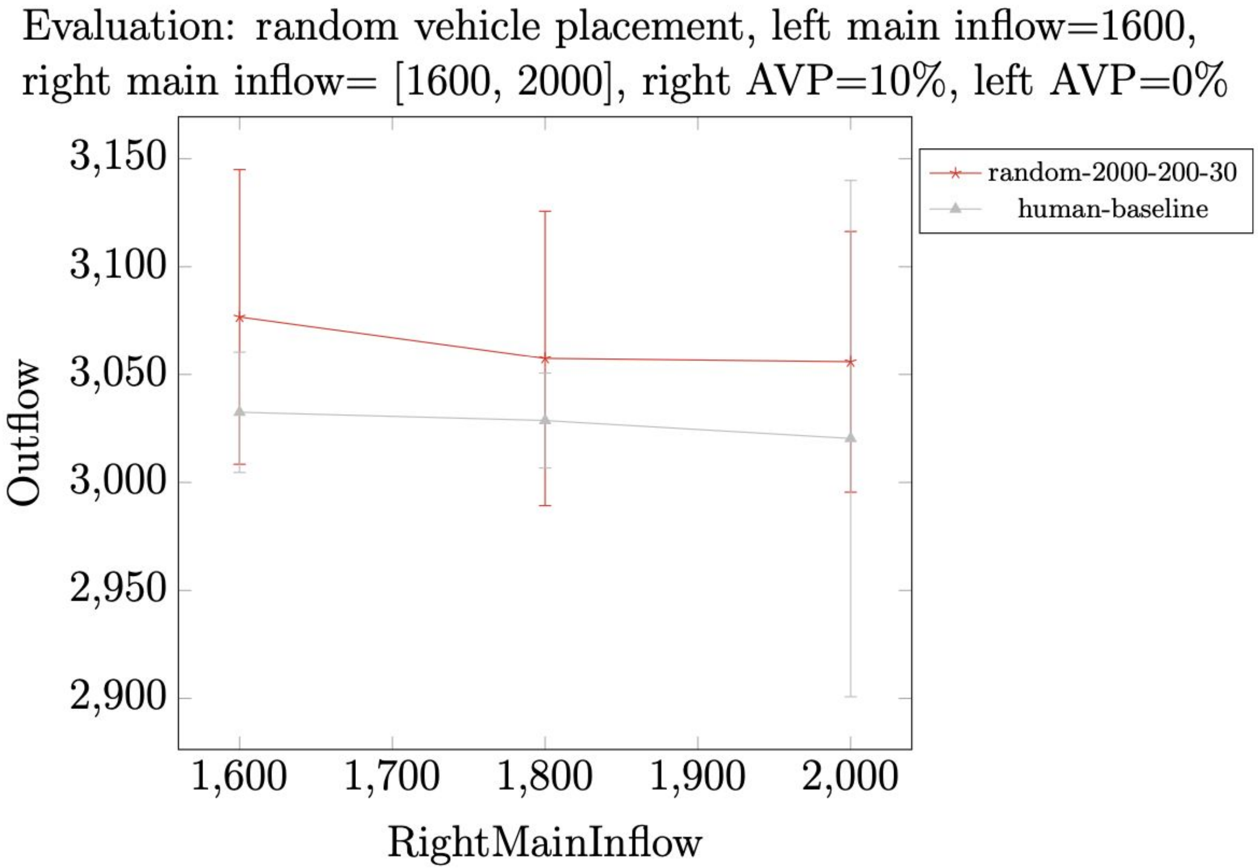

Conclusion and future work

- We have developed a single policy that is robust to:
  - Traffic flow
  - Fraction of AVs in traffic (AVP)
  - AV placement in traffic
  - Road geometry
    - Double ramps
    - Double lanes
- Limitations and Future work:
  - Existence of a left-lane policy in multi-lane scenarios.
  - All simulated human-driven vehicles share the same aggressiveness.
  - Generalize toward a wider variety of road geometries.
  - Still sim-2-real gap: noisy sensing, actuation delay

Learning a Robust Multiagent Driving Policy for Traffic Congestion Reduction

Yulin Zhang, William Macke, Jiaxun Cui, Daniel Urieli, Peter Stone

Learning Agent Research Group
The University of Texas at Austin
General Motors R&D Labs
Sony AI