Dissipating Traffic Jams with Piecewise Constant Control

Mayuri Sridhar and Cathy Wu (MIT)

NeurIPS Tackling Climate Change Workshop 2021

Background

- Carbon dioxide makes up 80% of US GHG emissions in 2019
- Transportation sector causes ~35% of total CO_2 emissions and 29% of GHG emissions [5]
- Previous studies show simple traffic management (speed management, shock wave suppression etc.) techniques can reduce emissions on the road by almost 20% [1]

Background

- Wu et al. shows that we can increase average speeds by up to 57% in idealized traffic settings with a small fraction of RL-controlled vehicles [6]
- Sridhar et al. shows that we can reproduce traffic smoothing results with piecewise constant policies (up to Δ=140) [3]
- This suggests we can use RL advice to empower human drivers to mitigate congestion for a near-term deployment

Overview

- We extend previous work to more difficult (and realistic) settings:
  - We show that piecewise constant policies (up to Δ=100) can dissipate shock waves even after the formation of a severe traffic jam
  - We show that piecewise constant policies are effective with local reward functions

Experimental Setup

- Single-lane ring road setup
  - 250m circumference
  - 22 vehicles
- Shown to develop traffic waves without the presence of AVs [4]
- Introduce a single agent trained with TRPO with reward function of their own velocity [2]

Traffic Jam Severity

Mild Traffic:

Heavy Traffic:

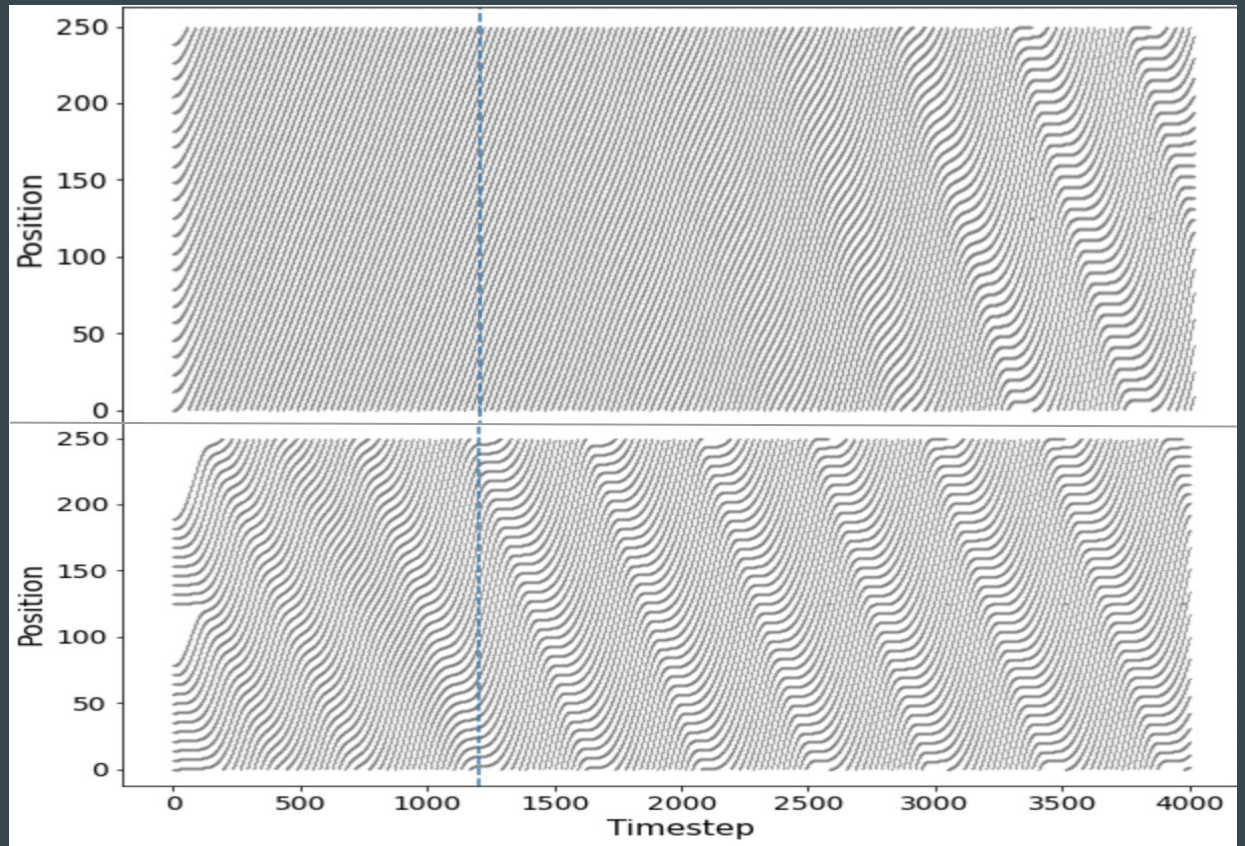

Traffic Mitigation

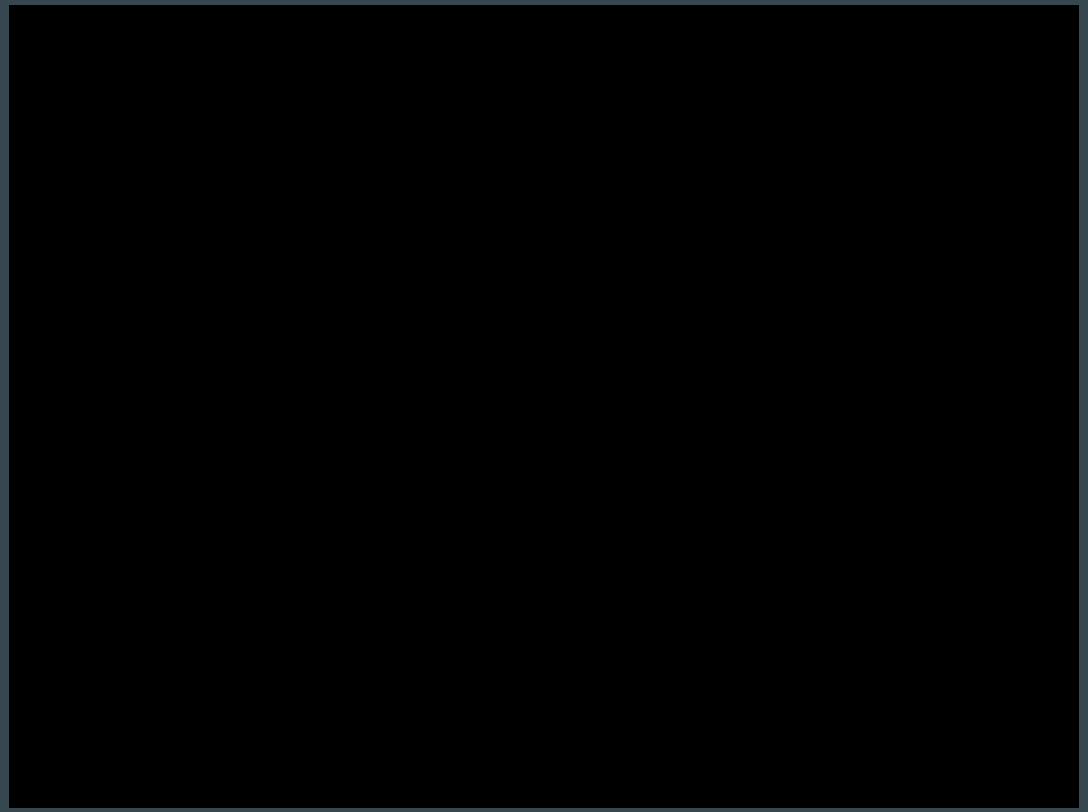

Results

- Congestion mitigation effects up to Δ=100
- Consistently beat human baseline velocity
- Lower velocity than Sridhar et al. due to traffic severity [3]

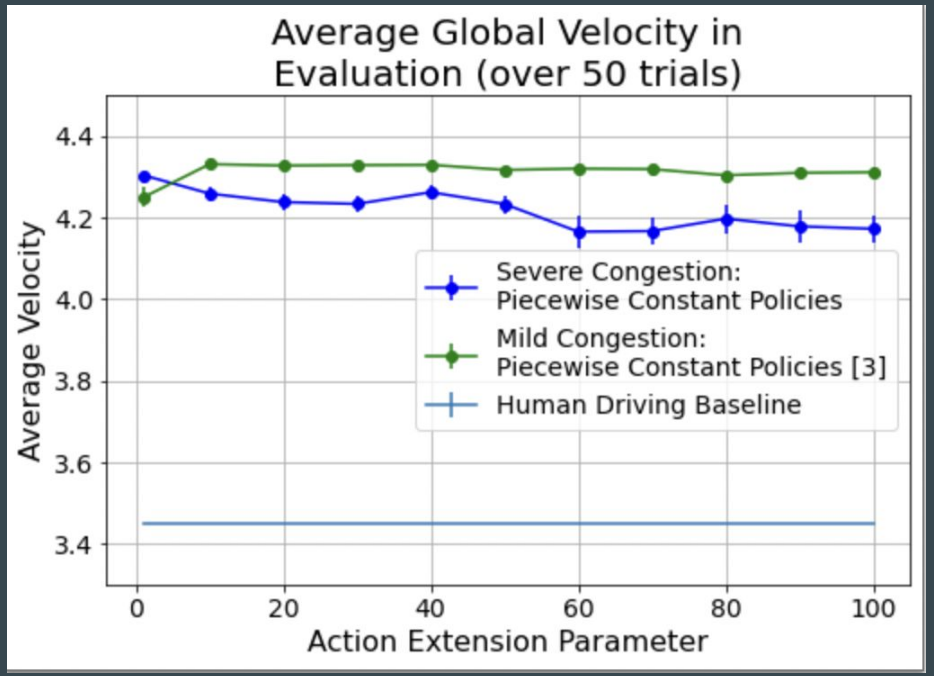

Contributions and next steps

- Our work suggests that piecewise constant policies work well in difficult settings, indicating potential for deployment
- We propose pursuing user studies in the future to see how well users can follow policies in practice to measure real-world congestion impacts

Acknowledgments

- Thank you to:
  - Zhongxia Yan and Cameron Hickert for advice on Flow framework and setup
  - Dr. Eric Horvitz for inspiring technical discussions
  - MIT Supercloud for computational resources
  - MIT-IBM Watson AI Lab for funding
  - Department of Defense for funding Mayuri through the National Defense Science and Engineering Graduate Fellowship

References

[1] BARTH, M., AND BORIBOONSOMSIN, K. Real-world carbon dioxide impacts of traffic congestion. Transportation Research Record 2058, 1 (2008), 163–171.

[2] SCHULMAN, J., LEVINE, S., ABBEEL, P., JORDAN, M., AND MORITZ, P. Trust region policy optimization. In Proceedings of the 32nd International Conference on Machine Learning (Lille, France, 07–09 Jul 2015), F. Bach and D. Blei, Eds., vol. 37 of Proceedings of Machine Learning Research, PMLR, pp. 1889–1897.

[3] SRIDHAR, M., AND WU, C. Piecewise constant policies for human-compatible congestion mitigation. In IEEE International Intelligent Transportation Systems Conference, 2021

[4] SUGIYAMA, Y., FUKUI, M., KIKUCHI, M., HASEBE, K., NAKAYAMA, A., NISHINARI, K., TADAKI, S.-I., AND YUKAWA, S. Traffic jams without bottlenecks - experimental evidence for the physical mechanism of the formation of a jam. New Journal of Physics 10 (03 2008), 33001.

[5] US ENVIRONMENTAL PROTECTION AGENCY. Inventory of U.S. greenhouse gas emissions and sinks: 1990-2019.

[6] WU, C., KREIDIEH, A., PARVATE, K., VINITSKY, E., AND BAYEN, A. Flow: A modular learning framework for mixed autonomy traffic. IEEE Transactions on Robotics (2021).