Low-Level Continuous Control for Highway Ramp Merging in Dense Traffic

Samuel Triest
University of Rochester, Computer Science Department

Adam Villaflor
Carnegie Mellon University, Robotics Institute

John M. Dolan
Carnegie Mellon University, Robotics Institute

Motivation

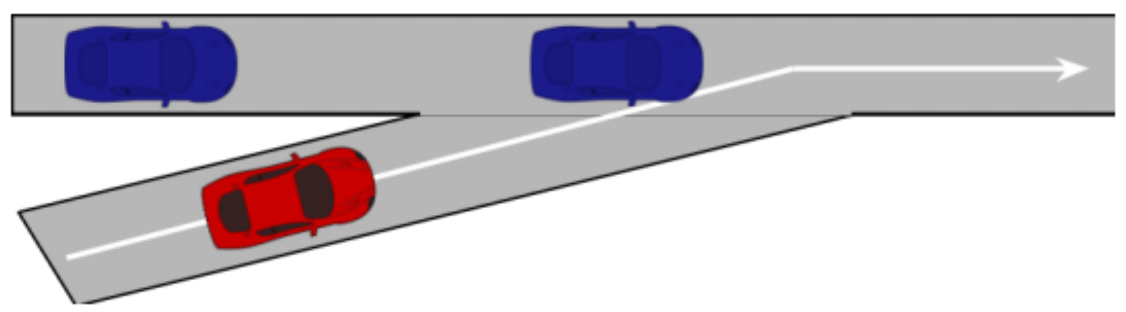

Fig. 1. The merging problem - the ego vehicle (in red) needs to move from the merging lane to the host lane without disrupting the host vehicles (in blue)

- A number of driving scenarios, including ramp merging, remain difficult for ADAS to handle.
  - High degree of interaction with human-driven vehicles
  - Must interact with drivers in a human-like way

Related Work

- Dong et al. use probabilistic graphical models to estimate intention, and ACC to generate merge trajectories.
- Kuefler, Bhattacharyya et al. use imitation learning to generate trajectories for general highway driving.
- Hu, Bouton et al. use reinforcement learning for merging to generate accelerations along a fixed merge path.
- Prior work assumes a path to a fixed merge point – no immediate application to certain ramp geometries, not necessarily human-like.

Generative Adversarial Imitation Learning with IDM Masking

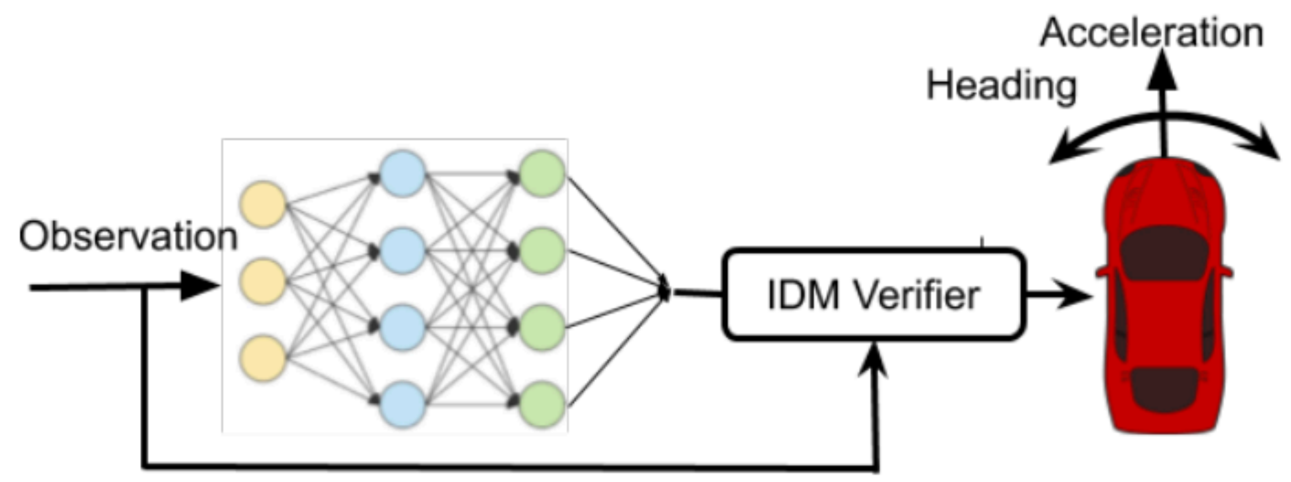

Fig. 2. The ego-vehicle is controlled using a deep neural network that controls the vehicle's heading and acceleration, given an observation. The policy output is verified using a distance-keeping model based on IDM.

- A policy is trained using GAIL on expert merging trajectories extracted from NGSIM.
  - Kinematic features are extracted from the ego-vehicle and 6 of tis neighbors
  - Low-level control outputs are used to generalize across ramp geometries, avoid fixed merge point
- IDM is used to provide safe bounds on the acceleration of the ego-vehicle
- Allows for verifiable low-level behavior

$$a_{IDM} = a_{max}(1 - (\frac{v_\alpha}{v_0})^4 - (\frac{s^*(v_\alpha, \Delta v_\alpha)}{s_\alpha})^2)$$

$$\text{where } s^*(v_\alpha, \Delta v_\alpha) = s_0 + v_\alpha T + \frac{v_\alpha \Delta v_\alpha}{2\sqrt{a_{max} b}}$$

$$\pi_{mask}(s) = \min(\pi(s), a_{IDM})$$

Results

	Slot-based	iPCB	MML-PGM	GAIL	GAIL + IDM
Success Rate	85.5%	84.2%	92.4%	62%	**91% (our work)**

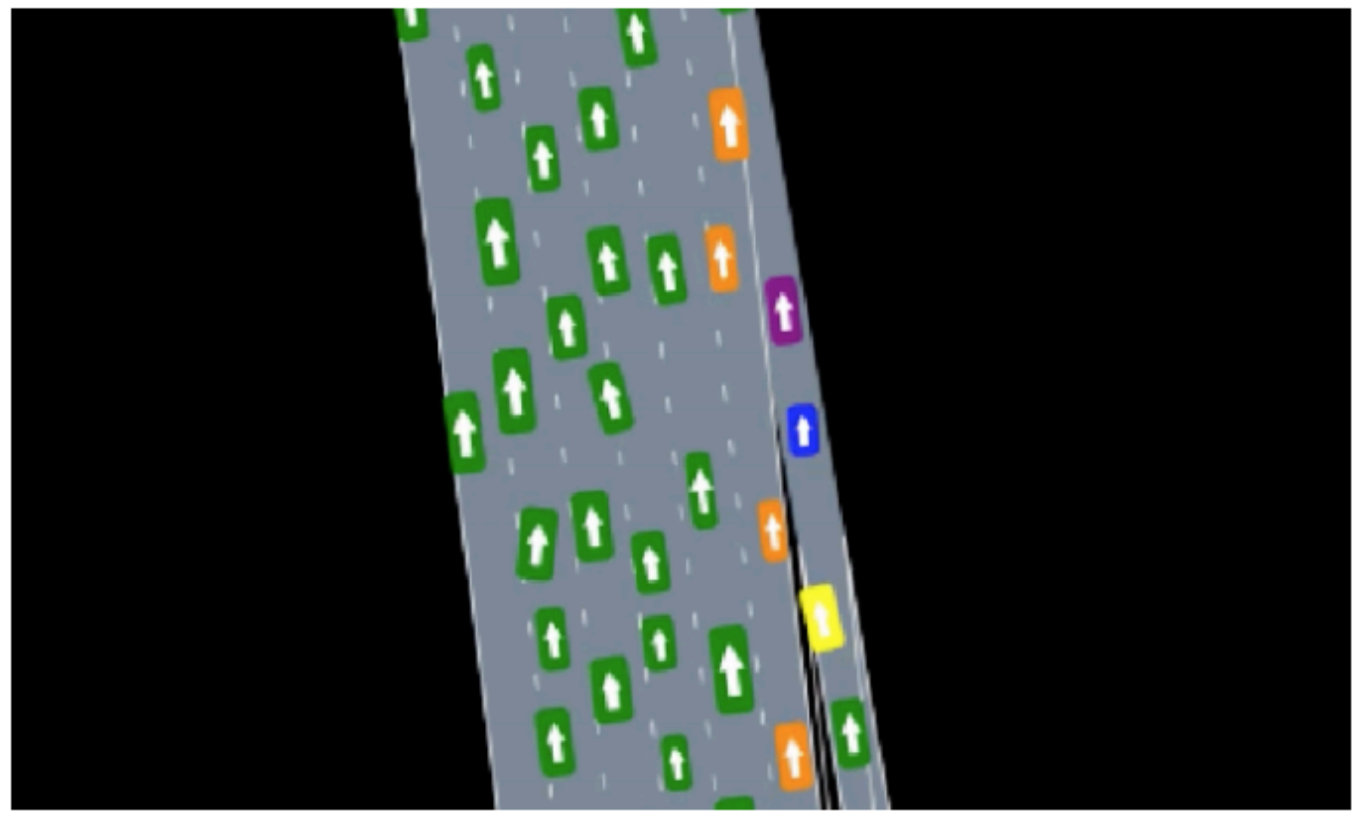

Fig. 3. Diagram of the vehicles that the ego-vehicle has access to - two in the the merging lane and four in the host lane. The ego-vehicle is in blue, the merging neighbors are in yellow, and the host neighbors are in orange. The leading vehicle that IDM is distance-keeping to is in purple. Green vehicles are unobserved by the ego-vehicle.

Discussion and Future Work

- GAIL with IDM masking performs about as well as other methods.
- The low-level control likely hurts GAIL performance without IDM.
- Will experiment with more sophisticated imitation learning methods (multi-agent, data augmenting) to get better results.

References

- C. Dong, J. M. Dolan, and B. Litkouhi, "Interactive ramp merging planning in autonomous driving: Multimerging leading pgm (mml-pgm)," in 2017 IEEE 20th International Conference on Intelligent Transportation Systems (ITSC). IEEE, 2017, pp. 1–6.
- Y. Hu, A. Nakhaei, M. Tomizuka, and K. Fujimura, "Interaction-aware decision making with adaptive strategies under merging scenarios," arXiv preprint arXiv:1904.06025, 2019
- M. Bouton, A. Nakhaei, K. Fujimura, and M. J. Kochenderfer, "Cooperation-aware reinforcement learning for merging in dense traffic," arXiv preprint arXiv:1906.11021, 2019
- A. Kuefler, J. Morton, T. Wheeler, and M. Kochenderfer, "Imitating driver behavior with generative adversarial networks," in 2017 IEEE Intelligent Vehicles Symposium (IV). IEEE, 2017, pp. 204–211

Acknowledgements

- Sam would like to thank:
  - The Robotics Institute and RISS
  - Professor John Dolan and Rachel Burcin and Adam Villaflor for their mentorship and advice
  - The NSF for providing funding