Multi-behavior Learning for Socially Compatible Autonomous Driving

Sanjula Jayawardana [†], Vindula Jayawardana [‡], Kaneeka Vidanage [†], Cathy Wu [‡]

[†] Informatics Institute of Technology [‡] MIT

Correspondence: sanjula.2019783@iit.ac.lk

IEEE ITSC 2023 BILBAO

Abstract

- Learning is used extensively in modeling emergent behaviors in autonomous driving.
- But, integrating autonomous vehicles (AVs) into the real world requires AVs to adhere to social norms of driving.
- However, current learning-based AV control methods often overlook this crucial aspect.
- We look at the problem of enabling socially compatible driving when AV control policies are learned.
- We show without social compatibility, AV policies tend to adopt dangerously competitive driving behaviors, while the incorporation of social compatibility fosters smoother vehicle maneuvers.

Motivation

Emergent behaviors in driving

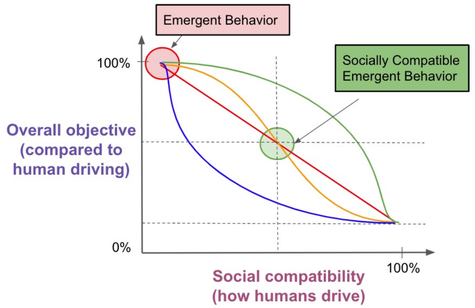

- Current works focus on learning emergent driving behaviors without considering the social compatibility of the learned behaviors.
- In the absence of social compatibility, the learned policies may lead to road rage and would face societal resistance.

Methodology

Learning human driving preferences

- Inverse reinforcement learning on real-world driving data.

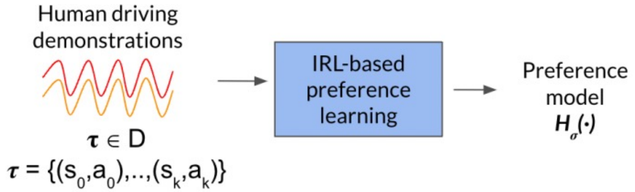

Modeling different driving behaviors

- Social value orientation to capture altruistic, prosocial, individualistic and competitive driving behaviors

$$g = cos(\phi) r_s(\cdot) + sin(\phi) r_o(\cdot)$$

self preference — preference of others

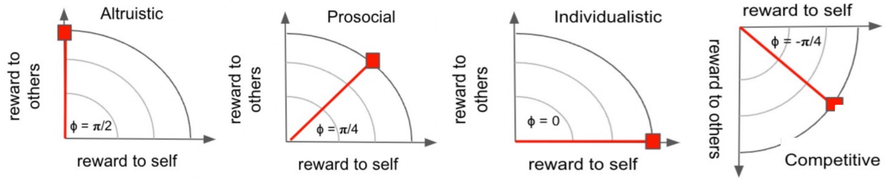

Learning multi-behavior socially compatible driving

- Multi-task learning to learn a unified policy that can demonstrate multiple driving behaviors.

objective: $\mathcal{J}(\pi) = \max_{\pi} \mathbb{E}\left[\sum_{\phi\in\Phi}\sum_{t=0}^{\infty} \gamma^t \Big[r(s_t,a_t) + \mu h(s_t,a_t)\Big] | s_0 = s, \pi(\cdot|\phi)\right]$

where $h(s_t,a_t) = H_\sigma^{AV}(s_t,a_t)\cos(\phi) + \Sigma_{i=1}^{m} q_i H_\sigma^i(s_t,a_t)\sin(\phi)$

Schematic overview

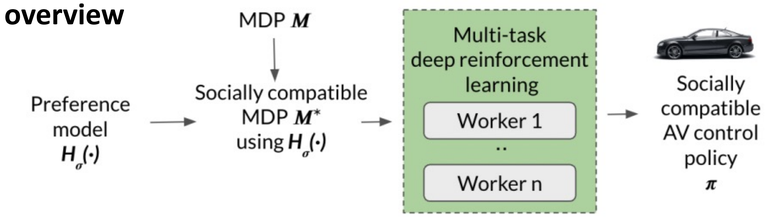

Results

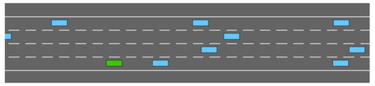

- Overall objective: Drive without collisions on a highway
- Simulation environment: Highway environment
- Baseline: Default DQN implementation (without social compatibility)

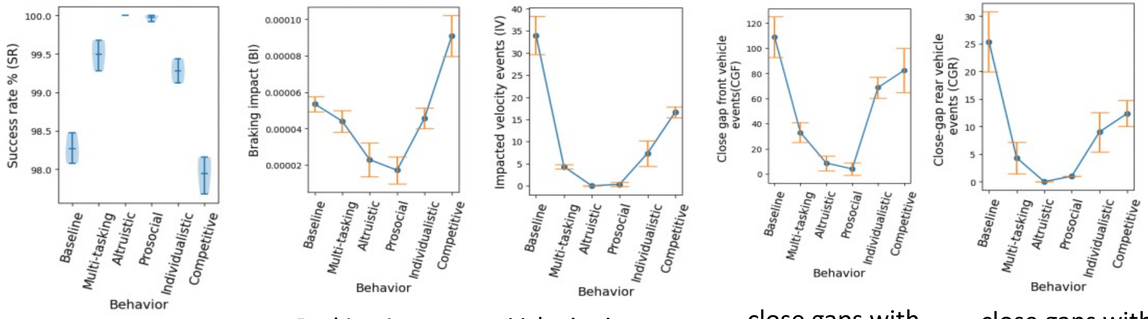

Success rate | Braking impact | Velocity impact | close gaps with front vehicle | close gaps with rear vehicle

- Without social compatibility, the default behavior (baseline) is significantly aggressive even more than the competitive behavior.
- Social compatibility can even help achieve the main objective.

Average human social preference and AV social preference →

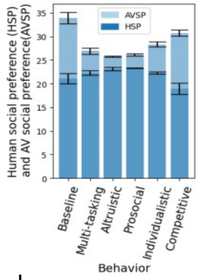

Conclusion

- Social compatibility is an important aspect when learning AV control policies.
- In the absence of social compatibility, acquired AV behaviors may take on an aggressive nature, potentially endangering those around the AVs.
- Integrating social compatibility into learning enables smooth driving and more human-like driving behaviors.