UCLA Samueli Computer Science

NEURAL INFORMATION PROCESSING SYSTEMS

PPL: Predictive Preference Learning from Human Interventions

Haoyuan Cai, Zhenghao Peng, Bolei Zhou

University of California, Los Angeles

NeurIPS 2025 **Spotlight**

Explore More

Interactive Imitation Learning from Human Interventions

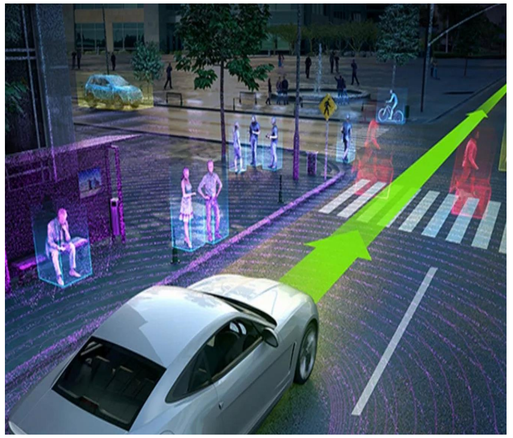

Autonomous driving

Robot grasping

Human intervenes and demonstrates correct actions **only when needed.**

Goal:

Reduce human effort and training time.

MetaDrive - At Training Time

Baseline: Proxy Value Propagation (PVP)

Peng, Mo, Duan, Li, Zhou. Learning from Active Human Involvement through Proxy Value Propagation. NeurIPS'23 Spotlight

Trajectory Prediction Model Guides Human Intervention

Ours: Predictive Preference Learning (PPL)

PPL Propagates Expert Interventions to Future Failures

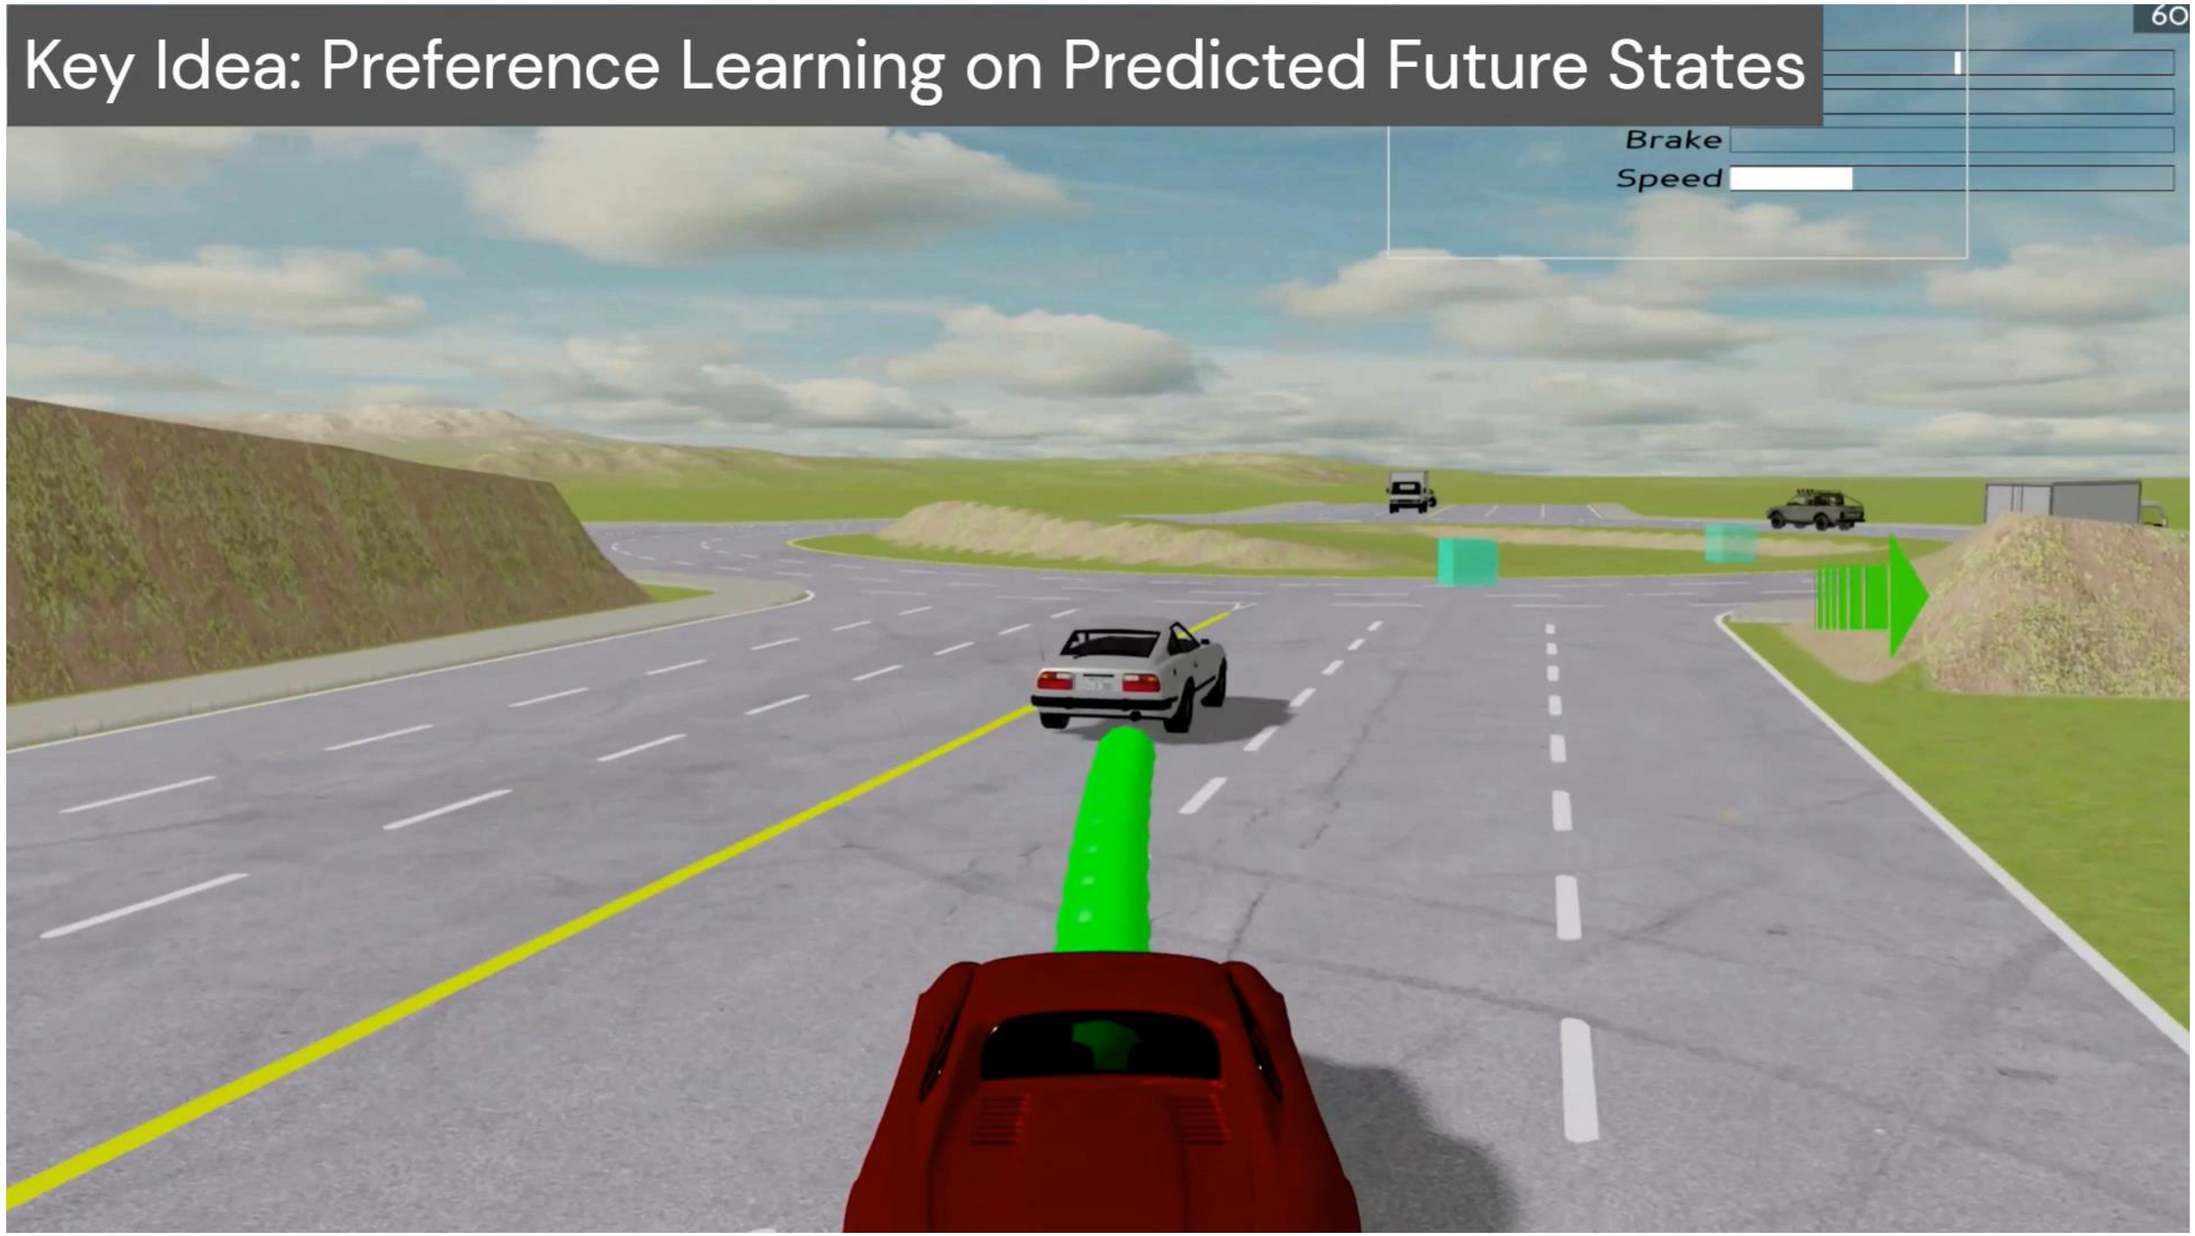

PPL trains a driving agent in 12 min & 1.8K demonstrations.

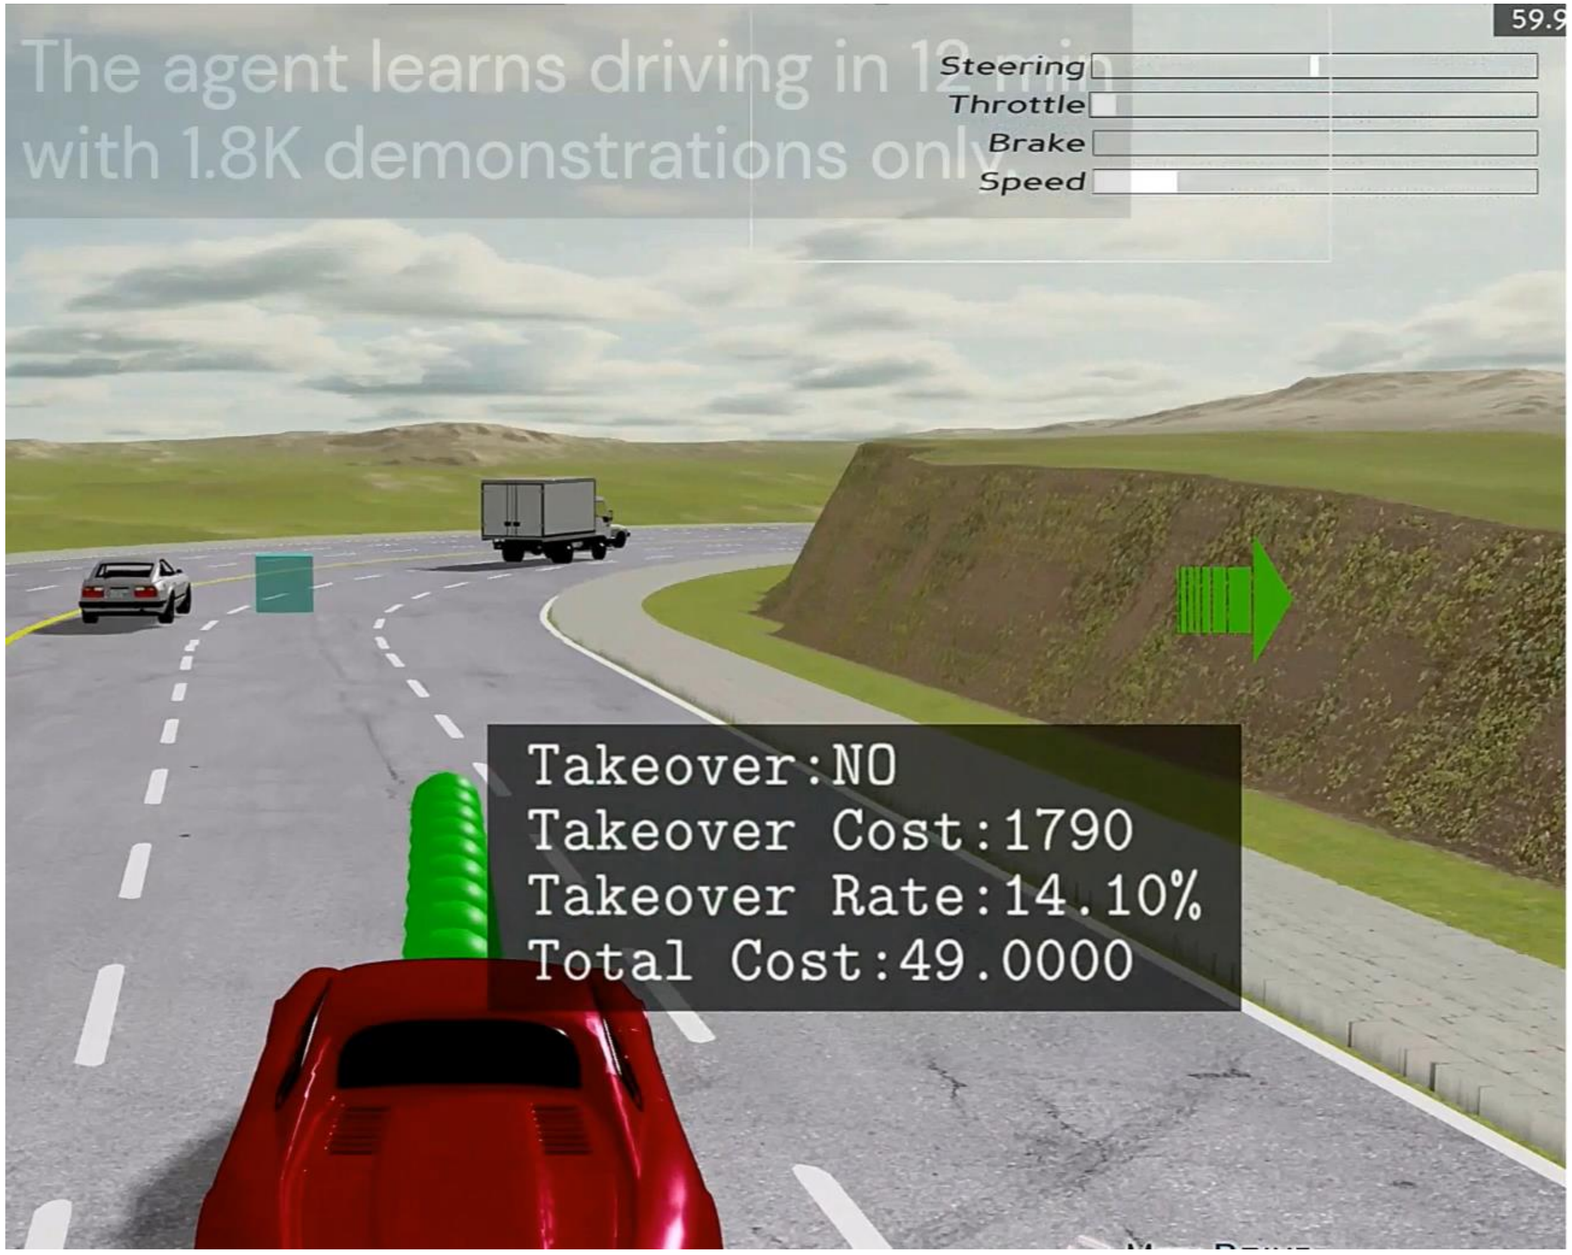

PPL Achieves 2x Improvement in Sample Efficiency

MetaDrive

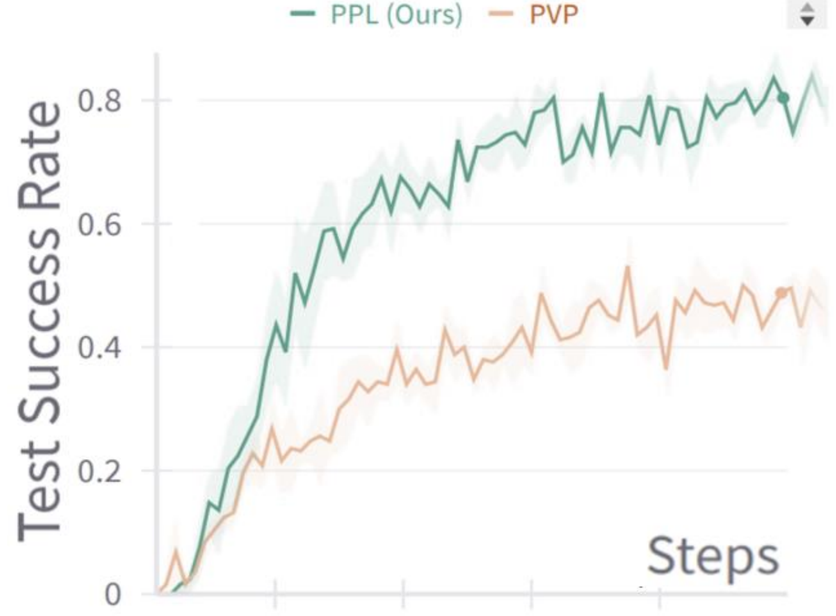

Table Wiping

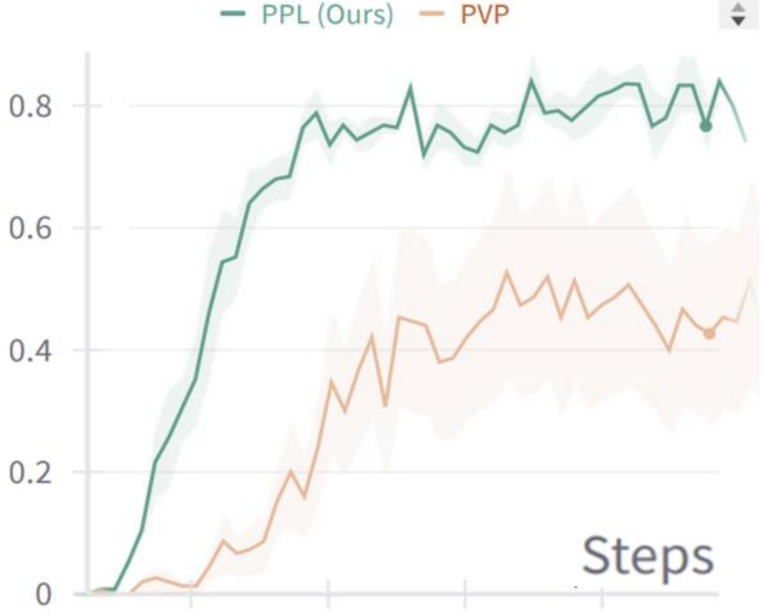

Nut Assembly

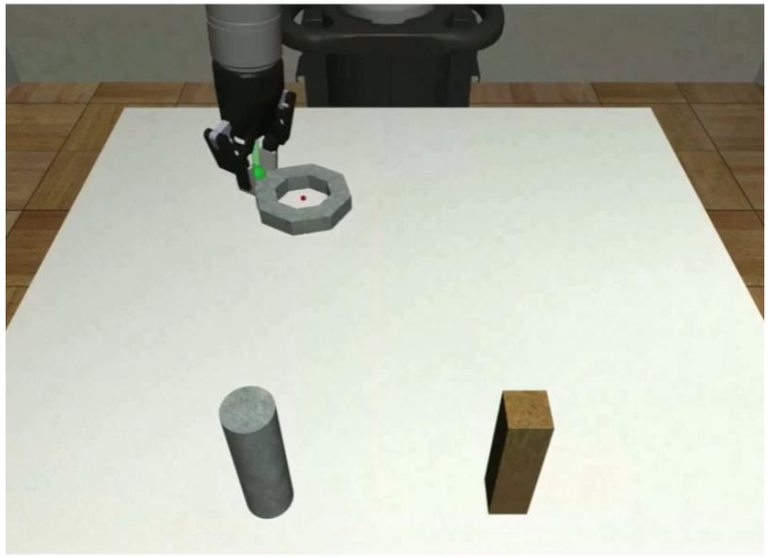

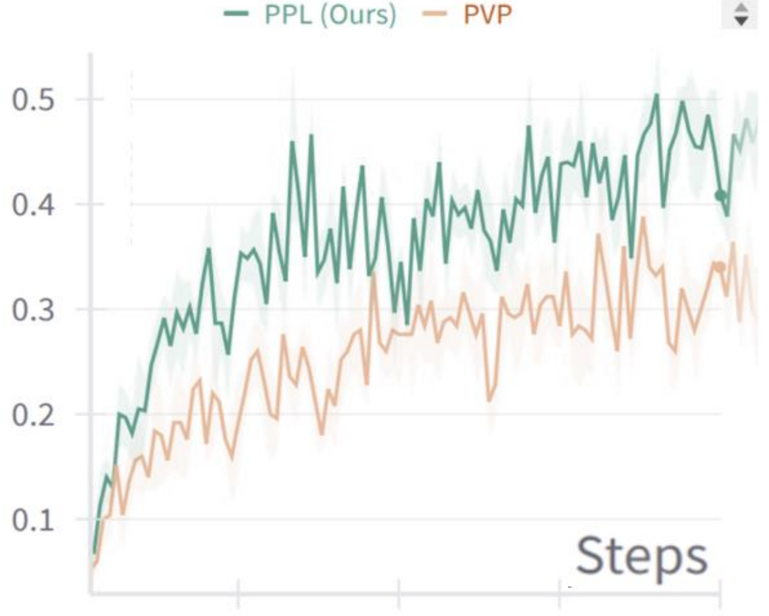

Predictive Preference Learning from Human Interventions

Haoyuan Cai, Zhenghao Peng, Bolei Zhou

NeurIPS 2025 **Spotlight**

1) Trajectory Prediction Facilitating Human Intervention

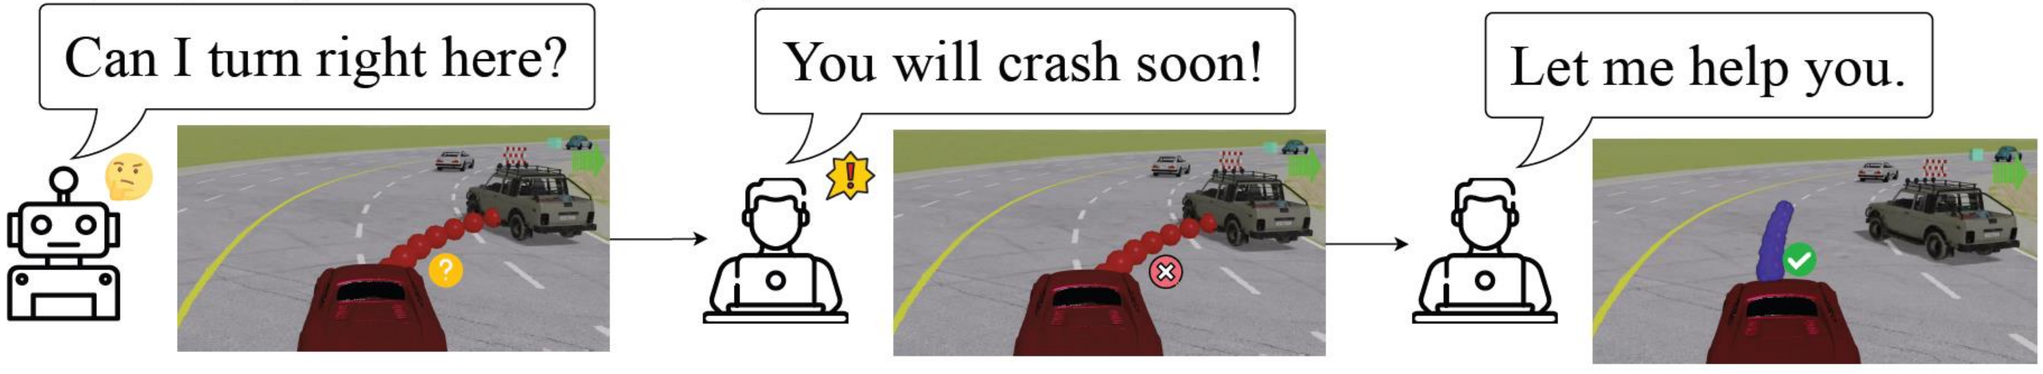

2) Learning Human Preferences in Forecasted Rollouts

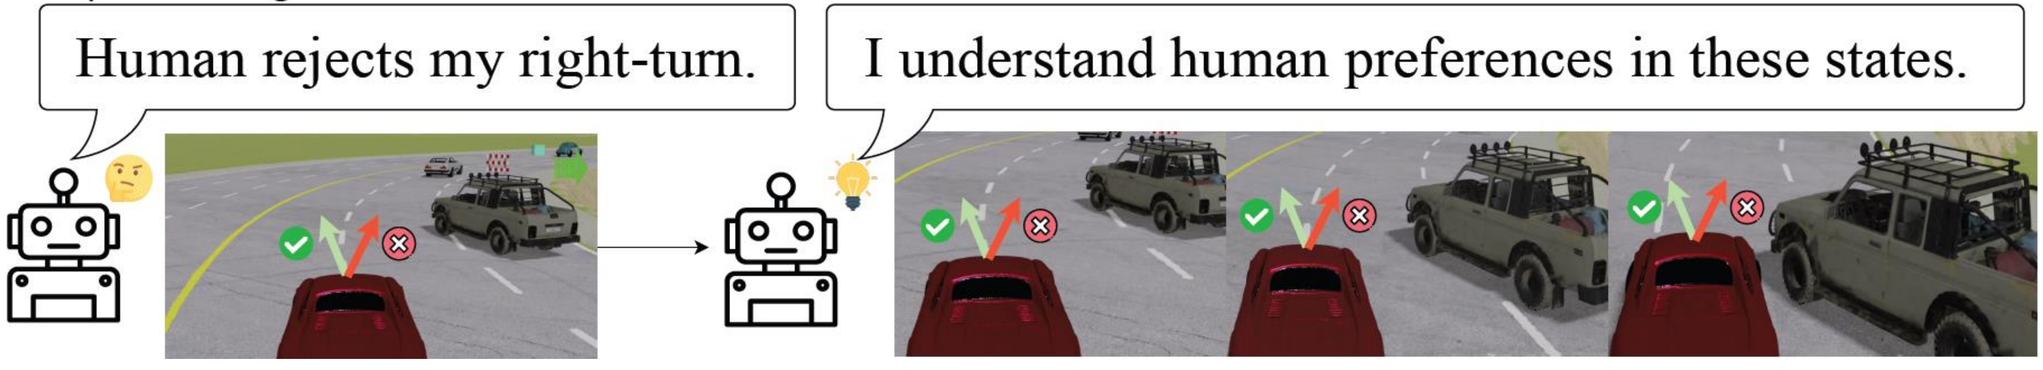

Explore More