

Safe Motion-Planning with Learning in the Loop

Emanuel Herberthson

Division of Automatic Control
Linköping University

May 20, 2026


- 1 Introduction
- 2 Problem
- 3 Results
- 4 Conclusions, Thoughts, and Future Work

This project looks to:

- Combine methods from both AI and optimal control into optimal motion planning
- Primarily reinforcement learning and generative AI
- Align with safety-critical applications. Thus achieve some level of guarantee


Optimal Motion Planning

- Planning the movement of a vehicle from A to B while minimizing some cost.
- Satisfy system dynamics and safety constraints.
- State lattice and motion primitives.
- Known costs to go between states, becomes a graph search.



Discrete State Lattice

(a) Discretization of the continuous state space into a lattice.



Motion Primitives

(b) Precomputed motion primitives connecting neighboring lattice states.

Problem

- Use reinforcement learning (RL) to learn motion primitives for a tractor with a dolly and a trailer. Start in an initial state, goal is to reach a given final state. Considering both forward and backward motion.
- Kinematic (bicycle) model from [1]. Assume $\alpha \in [-\frac{\pi}{6}, \frac{\pi}{6}]$.






Figure: An illustration of the tractor, dolly, and trailer system.

- Jackknifing when β_2, β_3 are too large.

- The system has an inherent symmetry, it can be mirrored in both x-axis and y-axis.
- The problem remains the same, except for that the optimal control changes sign.




(a) Mirroring in x-axis




(b) Mirroring in y-axis


Forward - slight turn



Forward - slight turn



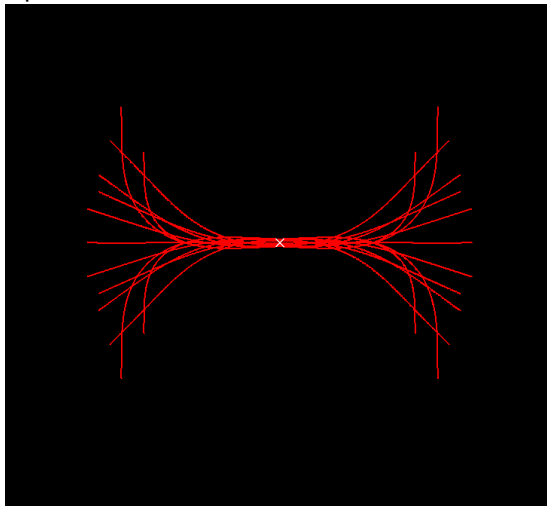
Forward - challenging turn



Forward - challenging turn



Forward motion primitives



Usage of symmetry - forward






Figure: Comparison of mirroring and no mirroring for forward motion

Reversing - slight turn



Reversing - slight turn




Reversing - challenging turn



Reversing - challenging turn



Backward motion primitives



Usage of symmetry - reverse

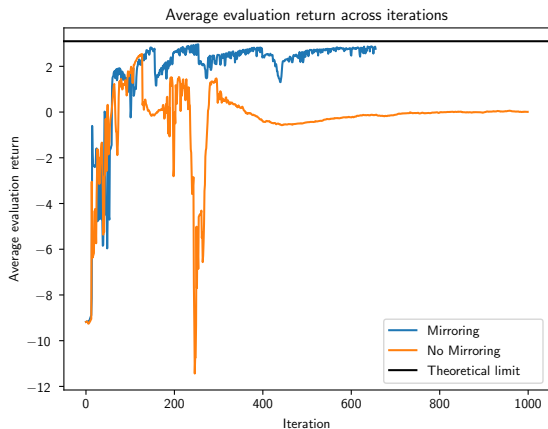


Figure: Comparison of mirroring and no mirroring for reverse movement

Usage of LQ-controller for stabilization - reverse

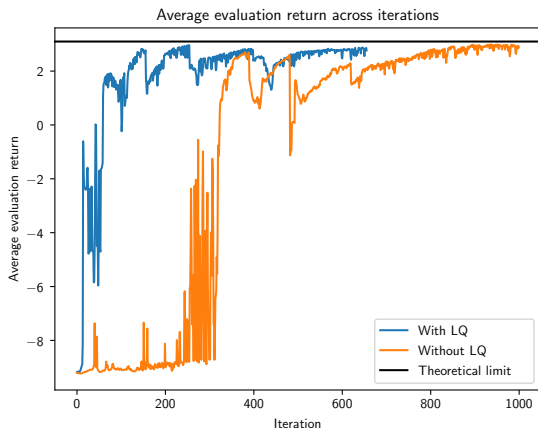


Figure: Comparison of using LQ-controller or not when reversing

- Usage of symmetry was crucial to enable learning.
- LQ-controller accelerated learning, but seemed to degrade performance. Could be due to the combination of residual policy learning and constrained actions.
- Better at sharp turns for forward motion. When reversing, perhaps must be on the edge of stability. In a noisy environment, this can be problematic.

- Convergence can still be improved. Verify, through optimal control, that the neural network can represent the policy.
- Improve final convergence through optimal control generated samples?
- The system is reversible, i.e. the controls for a given motion primitive is reversed if velocity changes sign. This should make it possible to accelerate learning further.
- How to combine residual policy learning with action constrained learning.

- [1] O. Ljungqvist, "Motion planning and feedback control techniques with applications to long tractor-trailer vehicles," Ph.D. dissertation, Linköping University, Department of Electrical Engineering, Automatic Control, Linköping, Sweden, 2020. DOI: 10.3384/diss.diva-165246. [Online]. Available: <https://www.diva-portal.org/smash/record.jsf?pid=diva2:1424832>.
- [2] W. Zhao, T. He, R. Chen, T. Wei, and C. Liu, *State-wise safe reinforcement learning: A survey*, 2023. arXiv: 2302.03122 [cs.LG]. [Online]. Available: <https://arxiv.org/abs/2302.03122>.
- [3] T. Silver, K. Allen, J. Tenenbaum, and L. P. Kaelbling, "Residual policy learning," *arXiv preprint*, 2018. arXiv: 1812.06298 [cs.RD]. [Online]. Available: <https://arxiv.org/abs/1812.06298>.

Thanks for listening!
Questions?