

Teaching a Race car to Drive Smarter: Blending Classical Control with Reinforcement Learning

Gokul M K, Arvind Easwaran

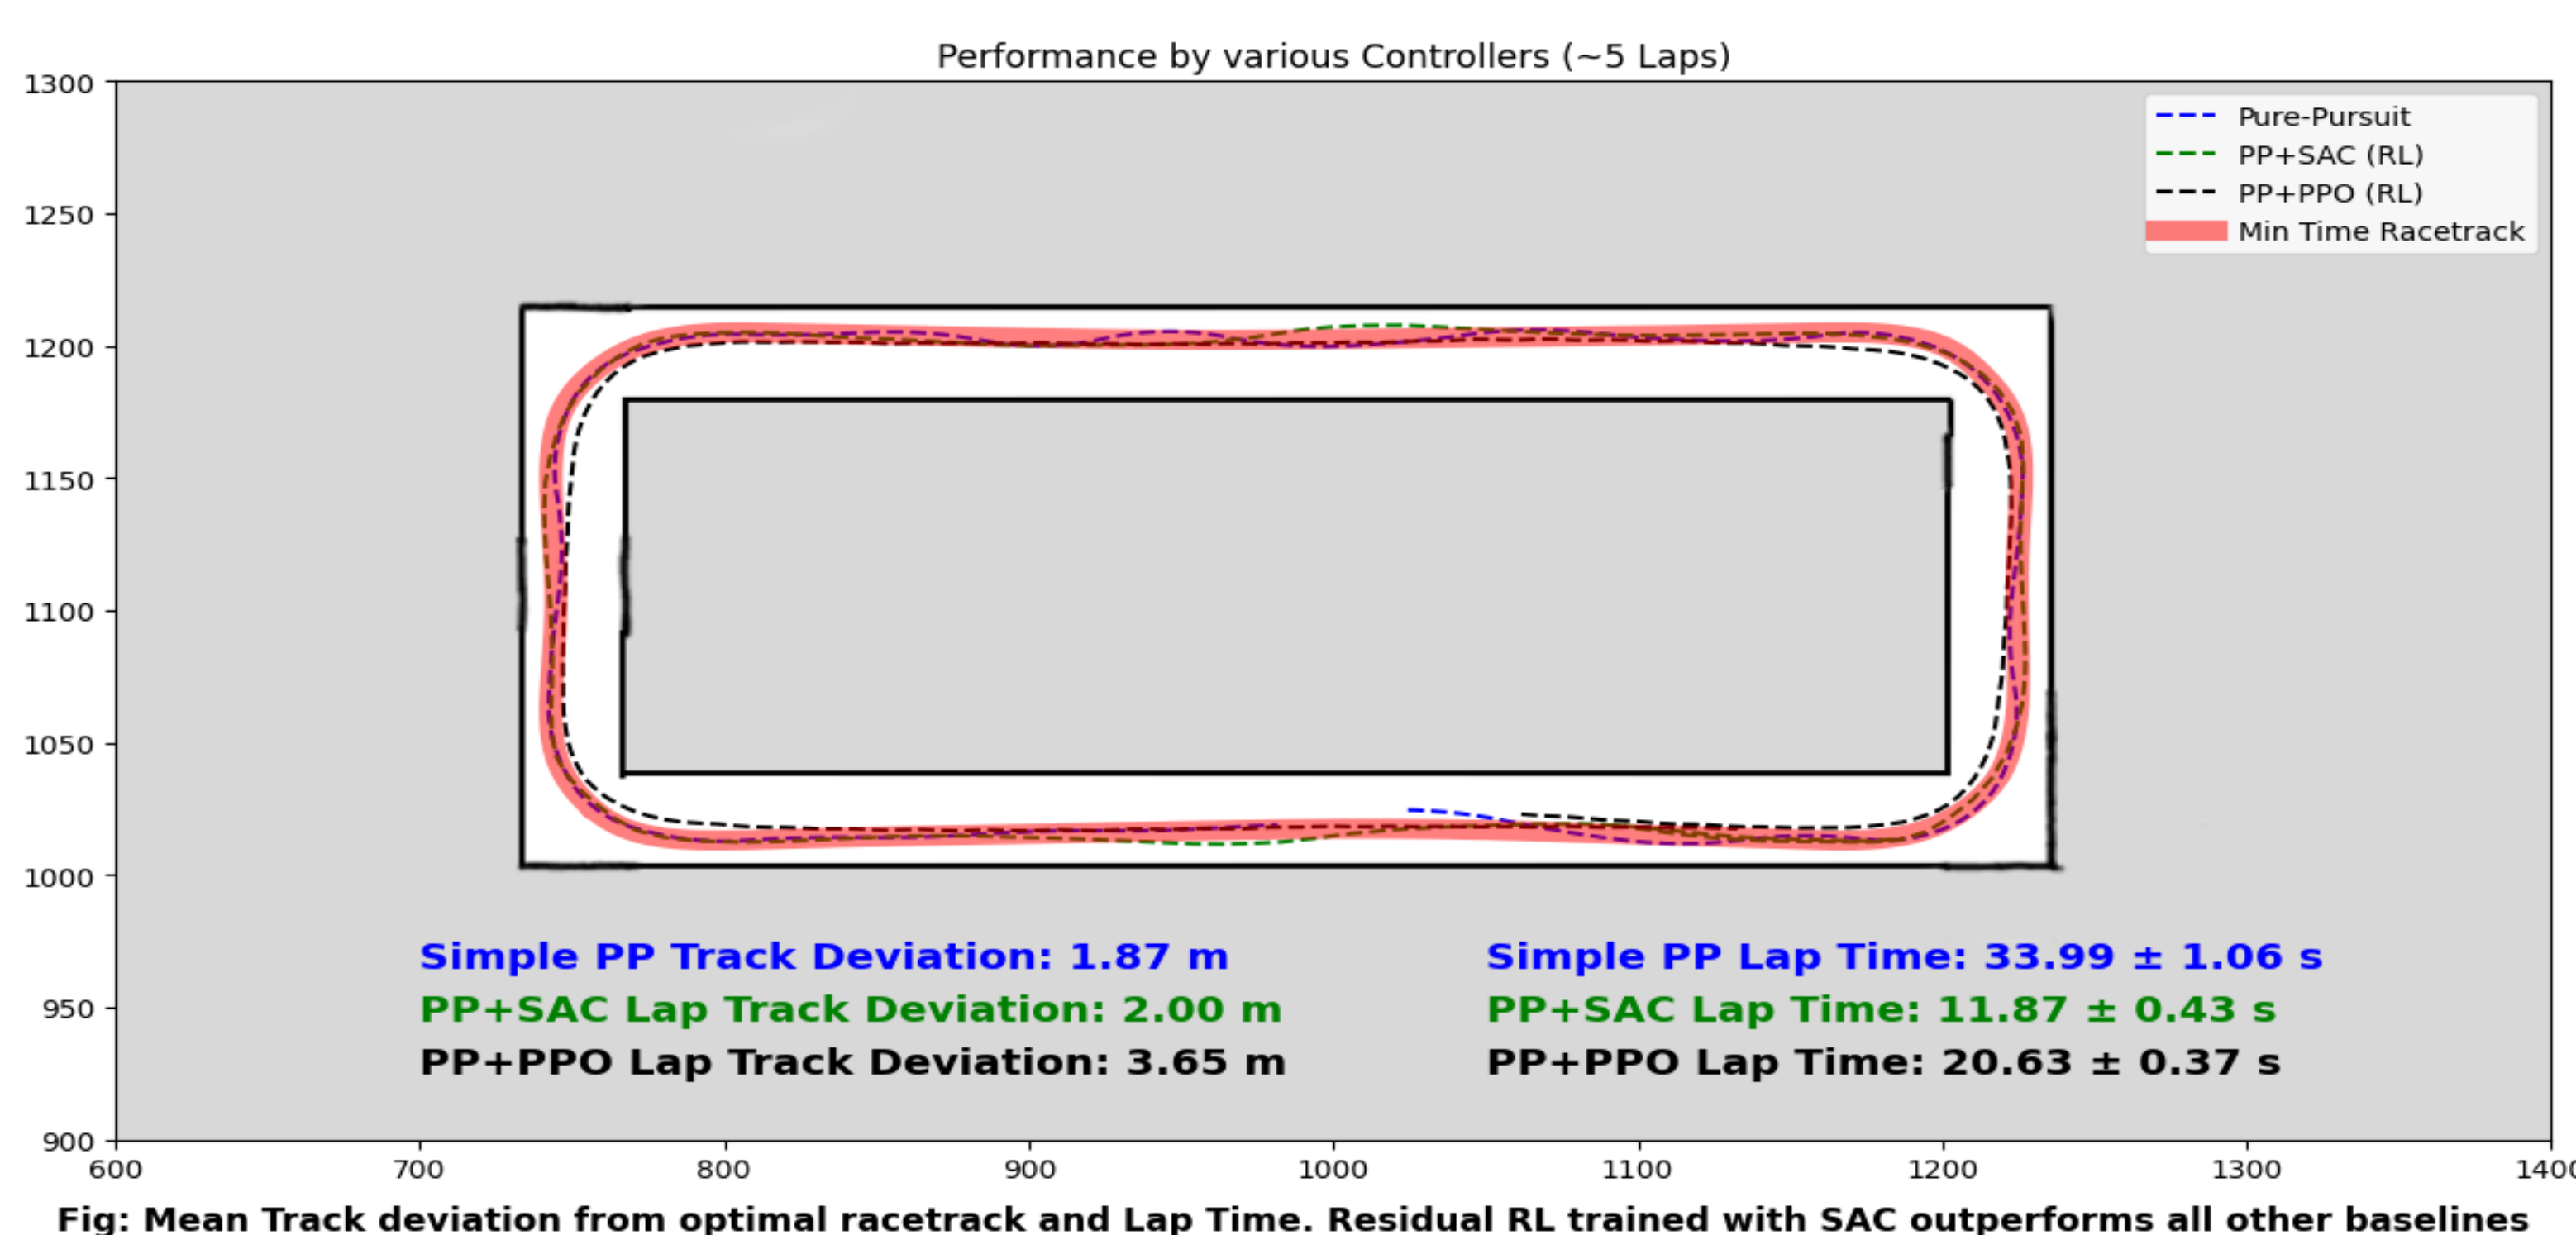

Hardware & Embedded Systems Lab, College of Computing & Data Science

Background:

- Classical Control (Pure Pursuit) is simple to implement, but **struggles** with high-speed maneuvers, dynamic environments and overtaking.
- It depends on many hyperparameters, requiring extensive **hand-tuning** for optimal performance, which is cumbersome and time-consuming.

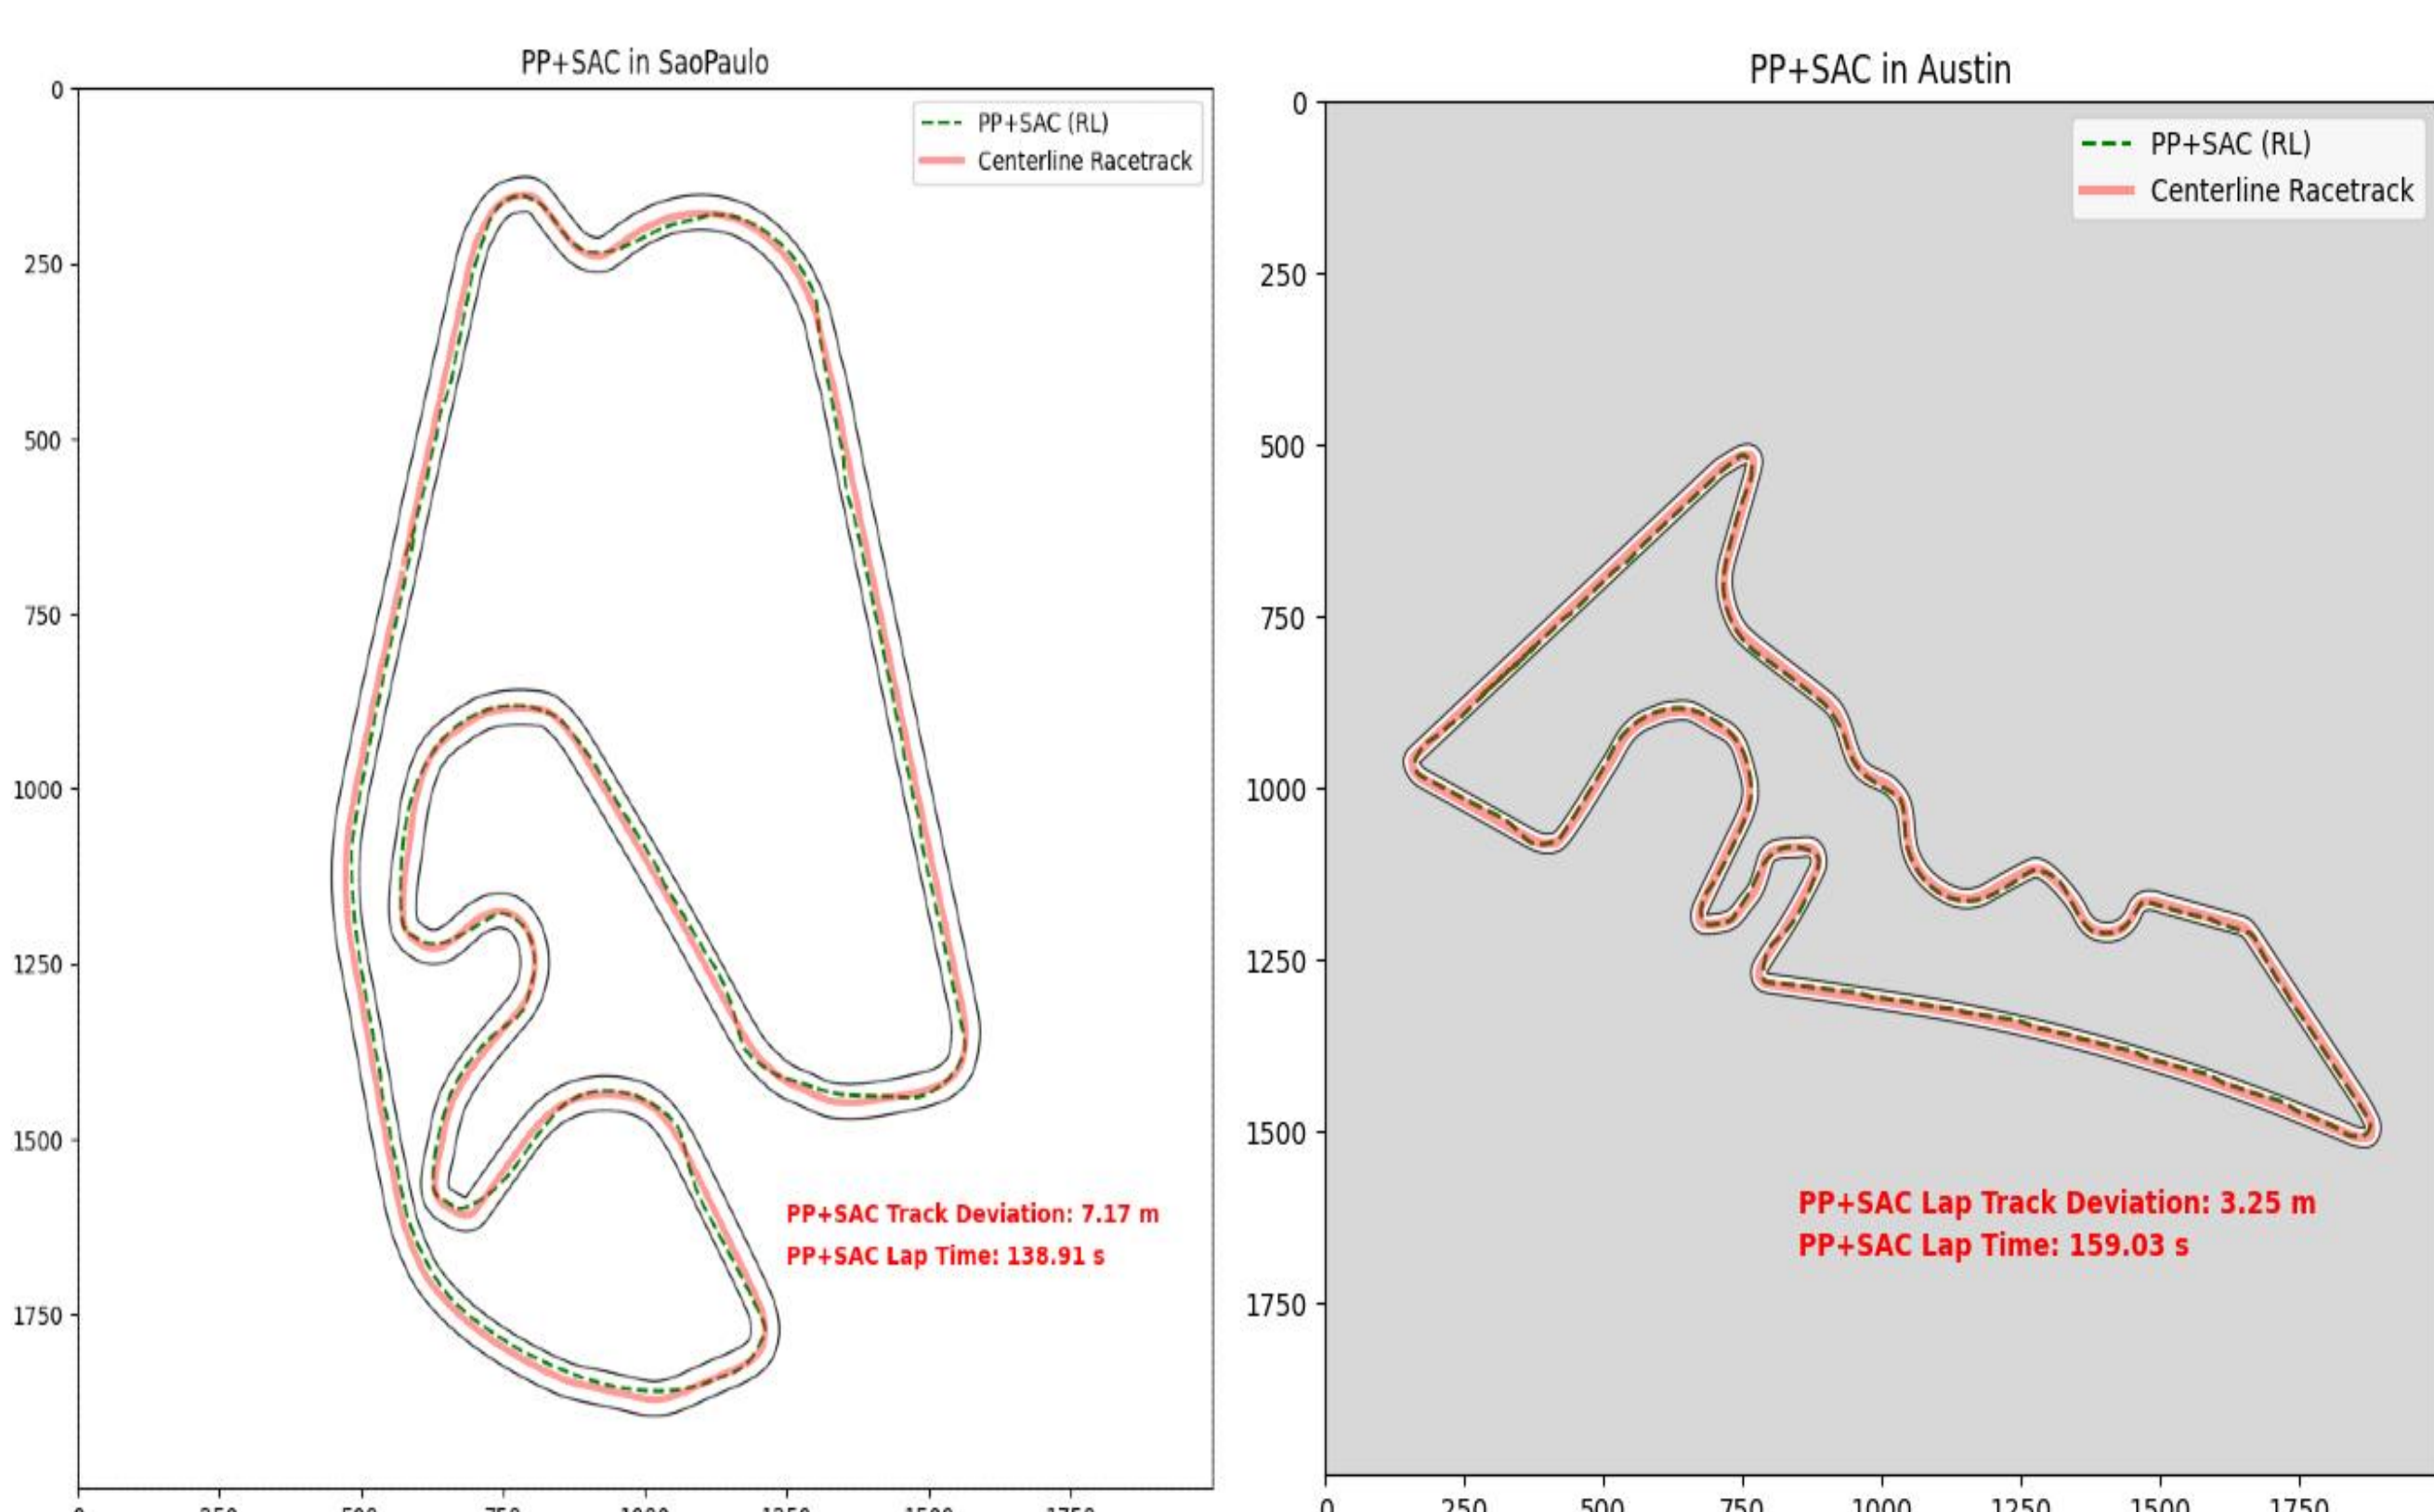

Aim:

- Improve the base controller to handle **unseen** and **challenging** racing scenarios using Reinforcement Learning.
- Use RL as a **residual** controller, so it learns from the base controller's mistakes and enhances speed, adaptability, and racing performance.
- Introduce **generalizability** and **adaptability** to the controller.



Observations:

- Residual RL significantly reduces the lap-times, demonstrating **improved** race car performance over the baseline.
- Among RL algorithms, **Soft-Actor Critic (SAC)** achieved the best results due to its **exploration-driven** learning.
- After training on a simple track, the agent successfully **generalized** to more complex environments.
- The agent learned to perform **overtaking** maneuvers even in the presence of dynamic obstacles.



Conclusion:

- Introducing RL as a **corrective** layer enhances the baseline's performance and improves generalization to complex scenarios.
- Future work:** Emphasize **Sim2Real** transfer by minimizing the gap between simulation and real-world deployment.

Acknowledgements:

- I am grateful to **GCF** and **Prof. Arvind Easwaran** in providing me this opportunity and to my buddies **Eduardo de Conto Fernandes**, **Subrat Prasad Panda**, **Michael Yuhas**, for their support.