Towards Safer Rehabilitation: Improving Gait Trajectory Tracking for Lower Limb Exoskeletons Using Offline Reinforcement Learning

M. Wong Sang[1], J. Narayan[2], B. Omarali[1], A. A. Faisal[1,2]

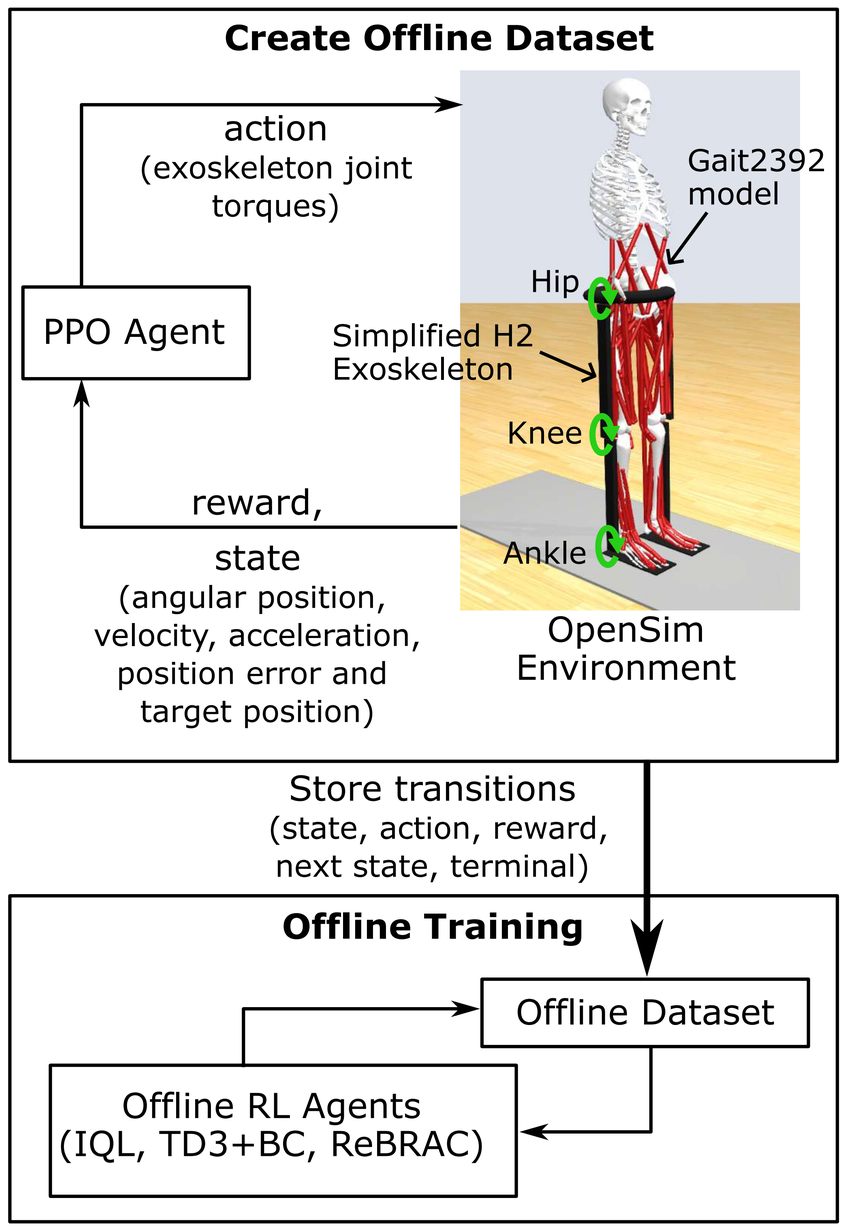
[1]Imperial College London,
[2]Universität Bayreuth
{a.faisal, b.omarali}@imperial.ac.uk

Introduction

Lower-limb exoskeletons are a scalable solution for gait rehabilitation, but traditional control methods struggle with user variability [1]. Reinforcement Learning (RL) offers adaptive control but poses safety concerns when applied online in clinical settings. Offline RL enables safe, data-driven learning from pre-collected experiences, avoiding real-time exploration risks [2]. This study presents the first application of offline RL for multi-joint gait trajectory tracking in lower-limb exoskeletons.

Methodology

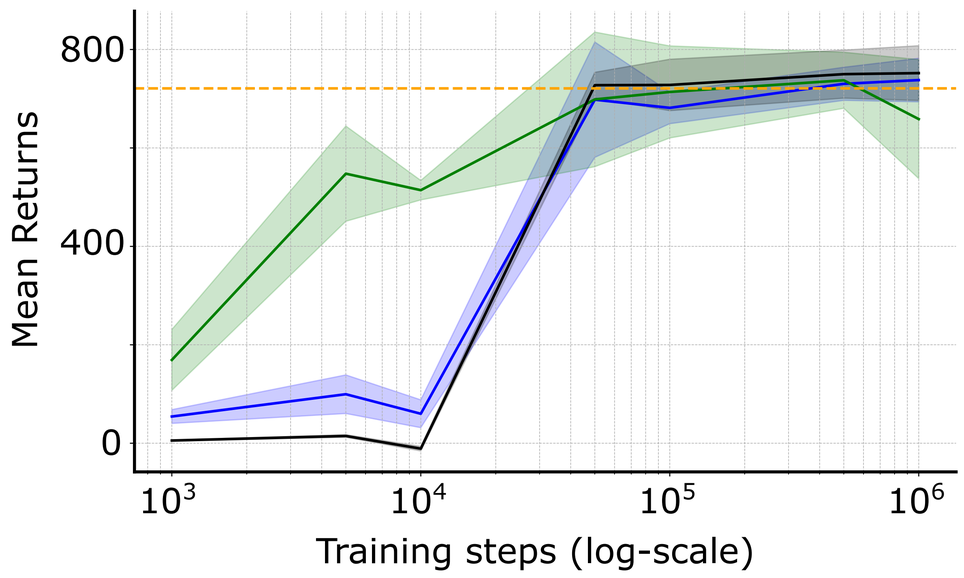
We simulate lower-limb gait using the OpenSim Gait2392 model paired with a simplified H2 exoskeleton [3] featuring actuators at the hip, knee, and ankle. A reward function promotes accurate trajectory tracking and joint safety. An offline dataset of 207 PPO-generated episodes (≈200k transitions) is used to train three offline RL algorithms: IQL, TD3+BC, and ReBRAC. Algorithms were implemented using Stable-Baselines3 [4] and d3rlpy, with hyperparameters tuned via Optuna to optimize learning performance.

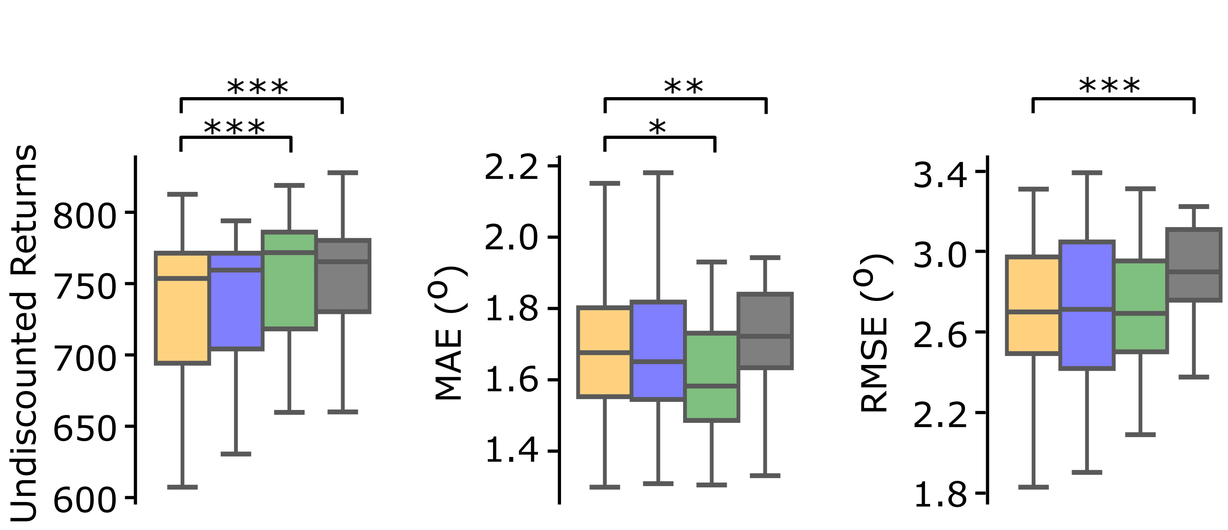


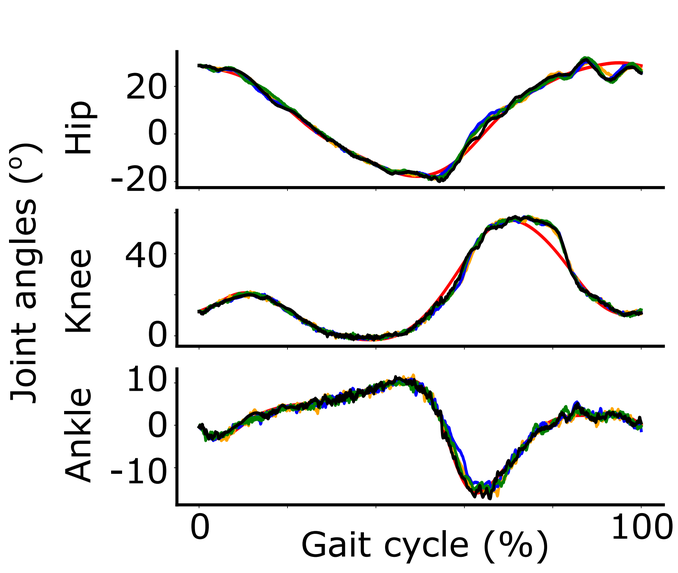
Results

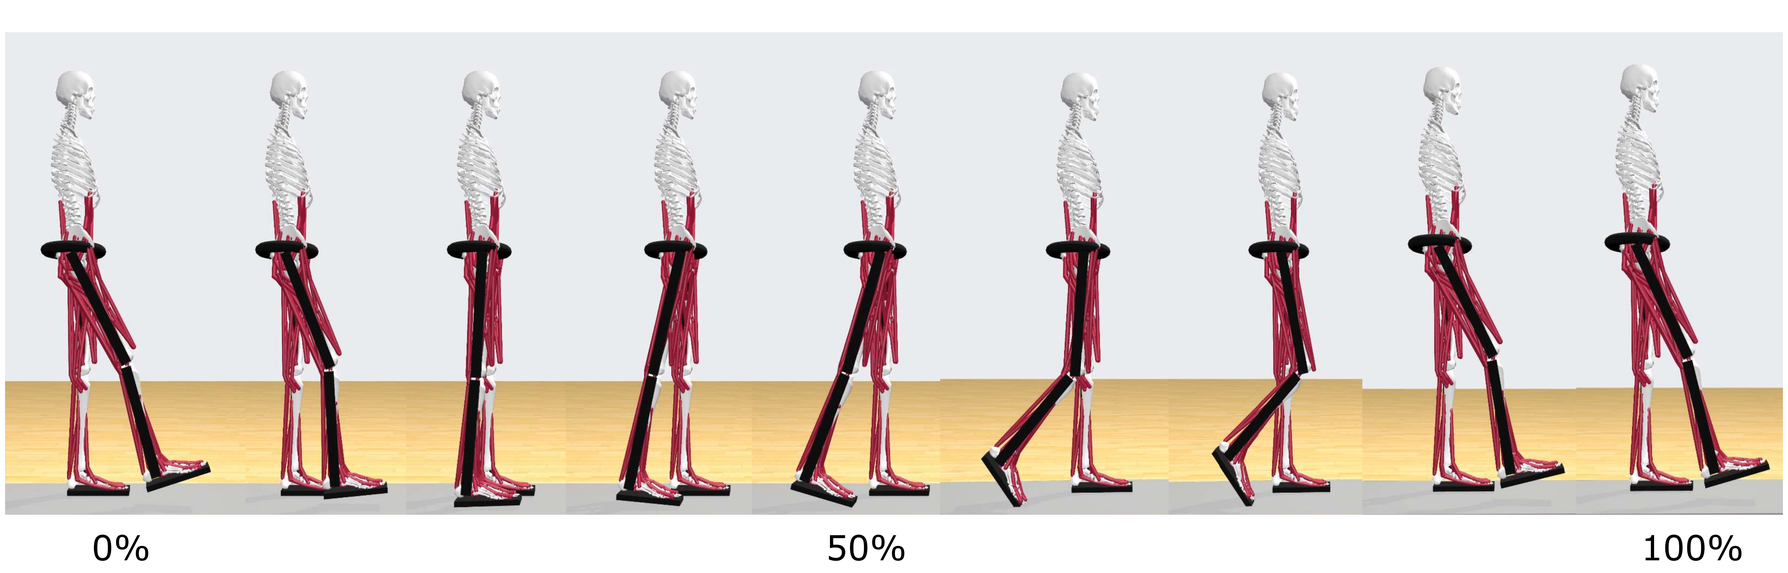
Models converged within the first 50–100k steps, highlighting the efficiency and robustness of offline RL in gait control tasks. TD3+BC showed the best balance between trajectory tracking accuracy and return, while ReBRAC achieved the highest returns with slightly worse tracking. IQL underperformed, showing no significant improvement over the original PPO data.

Model	Hip (°)		Knee (°)		Ankle (°)	
	MAE	RMSE	MAE	RMSE	MAE	RMSE
PPO	1.76±0.23	2.63±0.39	2.05±0.19	3.31±0.35	1.22±0.10	1.97±0.30
IQL	1.64±0.13	2.37±0.22	2.13±0.30	3.55±0.58	1.24±0.13	2.05±0.35
TD3+BC	**1.59±0.16**	**2.38±0.23**	**2.13±0.18**	**3.64±0.31**	**1.06±0.12**	**1.79±0.31**
ReBRAC	1.69±0.20	2.54±0.25	2.25±0.17	3.78±0.33	1.23±0.07	2.13±0.22









Conclusion

This work demonstrates that offline RL offers a safe and effective solution for trajectory control of lower-limb exoskeletons. By enabling robust multi-joint gait tracking entirely in simulation, our approach lays the foundation for clinical deployment. Future directions include applying offline RL to adaptive control of lower-limb exoskeletons in active-assist mode, expanding dataset diversity, and exploring sim-to-real transfer.

References

[1] B. Kalita, J. Narayan, and S. K. Dwivedy, "Development of active lower limb robotic-based orthosis and exoskeleton devices: a systematic review," Int. J. Soc. Rob., 2021
[2] M. Li, X. Gao, Y. Wen, J. Si, and H. H. Huang, "Offline policy iteration based reinforcement learning controller for online robotic knee prosthesis parameter tuning," ICRA2019
[3] M. Bortole, A. Venkatakrishnan, F. Zhu, J. C. Moreno, G.E. Francisco, J. L. Pons, J. L. Contreras-Vidal "The H2 robotic exoskeleton for gait rehabilitation after stroke: early findings from a clinical study." J. Neuroeng. and rehab 2015
[4] A. Raffin, A. Hill, A. Gleave, A. Kanervisto, M. Ernestus, and N. Dormann "Stable-baselines3: Reliable reinforcement learning implementations." JMLR 2021.