Learning Backswing Kicks with Deep Reinforcement Learning

Nico Bohlinger, Klaus Dorer

Hochschule Offenburg, Germany, contact: klaus.dorer@hs-offenburg.de

Deep Reinforcement Learning algorithms are used by the teams in the 3D simulation league since Abreu et al. [1] showed how to learn a running behavior with Proximal Policy Optimization (PPO). Later work [2] applied this approach to the task of kicking a ball as far and precise as possible.

While the kicks learned in this way with PPO were better in regards to the kickable area, the avoidance of falls and general reliability, they could not achieve the same level of kick distance as kicks learned with genetic algorithms. The reason for this discrepancy in kick distance is a result of the agent never learning to swing its leg back and forward to hit the ball with more power. Finding a movement like that in the search space of the agent would result in initially bad rewards for kicking the ball as far as possible due to the instability of this movement if not learned properly yet.

To guide the agent into the right regime in the search space we apply a curriculum consisting of an initial phase with a specific reward to only learn the needed backswing and follow-through movement. Then in a second and third learning phase we base the reward on the actual kick distance and fine-tune on variable kick distances and directions.

The improvement in gameplay was shown in a series of 200 games of two identical teams of eleven robots where the team using the new backswing kick scored an average of 4.375 goals compared to 0.535 goals of the team using our previous set of kicks consisting of a mix of non-backswing kicks learned with PPO and genetic kicks learned with CMA-ES [3].

Due to its high accuracy, flexibility and better performance the new backswing kick can completely replace the previous set of kicks. This means that only one single model has to be trained per robot type to cover all necessary kick distances and directions that are needed during a game.

References

1. Abreu M, Lau N, Sousa A and Reis L P (2019). Learning low level skills from scratch for humanoid robot soccer using deep reinforcement learning. In Proceedings of 2019 International Conference on Autonomous Robot Systems and Competitions (ICARSC). IEEE.
2. Spitznagel, Martin & Weiler, David & Dorer, Klaus. (2021). Deep Reinforcement Multi-Directional Kick-Learning of a Simulated Robot with Toes. 104-110. 10.1109/ICARSC52212.2021.9429811.
3. Dorer K (2017). Learning to Use Toes in a Humanoid Robot. In Akiyama H, Obst O, Sammut C, Tonidandel F: RoboCup 2017: Robot World Cup XXI, Nagoya, Japan, Springer Verlag.

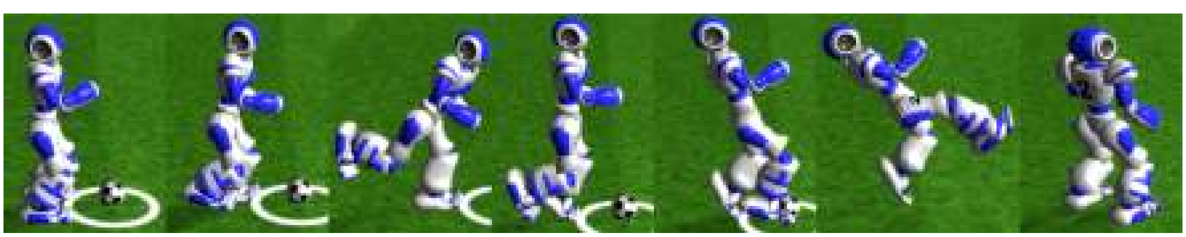

Fig. 1. Visualization of the learned backswing and follow-through phases.