Deep Reinforcement Keep Learning

Simon Rösch, Klaus Dorer

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

Over the last years, RoboCup Teams have been taking advantage of Deep Reinforcement Learning (DRL) to improve various aspects of their performance. These improvements range between running at greater speeds than was possible before and placing a shot in a specific area of the goal. To combat one of these improvements, namely placing a shot in a specific area of the goal, we have decided to use DRL to improve our goalkeepers performance for such hard-to-keep shots.

Trying to let the agent learn goalkeeping on all possible shots leads to it mostly learning to keep the easiest shots. These don't require it to time its movement exactly with the moment the ball passes. Therefore, the agent only really blocks shots for which it has to stay upright and move a few steps or lay down on the ground and wait for the ball to hit it. These are the same shots that it can already keep without DRL.

To combat this we isolate the hard-to-keep shots and only use these as our training dataset. We also include a timing-factor in our reward-function to guide the agent to time its movement. Using these methods the agent learns a movement consisting of a sidestep and extending its arms towards the point where the ball is going to pass it in the right moment. On one hundred kicks our previous goalkeeper was only able to keep 5 of these hard-to-keep shots. Our new goalkeeper is able to keep 36 out of a hundred of them.

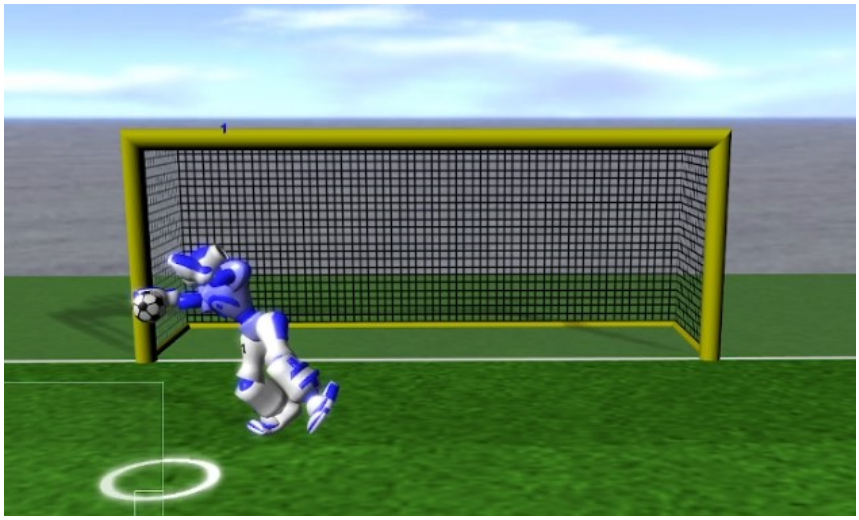

Fig. 1. Goalie in a keep movement.