Machine Learning for Robotics SS 2017
Instructor: Georg Martius <georg.martius@tuebingen.mpg.de>
Final Project
Due-date: 03.08.2017? Exact date will be fixed soon.
Intermediate steps discussed on 13 and 20.07.
Version 2, July 12, 2017

1 Reinforcement Learning Project

The final project of the lecture will be on Reinforcement learning. You can form teams of 2. You will implement several reinforcement learning algorithms and evaluate their performance on different scenarios. We are using Open AI gym, see `https://gym.openai.com/`. For the code and installation instructions see `https://github.com/openai/gym`.
All teams will compete against each other on the same environments. There are intermediate checkpoints, see below, that we will discuss in the remaining recitation sessions.
For the final evaluation, you have to prepare:

(a) A report with:

- description of the chosen methods and your potential modifications,
- experimental evaluation for all the methods and environments, and a
- discussion.

Each team member should have an independent contribution that is clearly marked in the report.

(b) A presentation $\sim$20 min

(c) The source code

Everything needs to be finished/presented at the "exam" day (probably 03.08.2017).

1.1 Environments

You should solve the following environments:

(a) Pendulum-v0

(b) LunarLander-v2

(c) BipedalWalker-v2

(d) BipedalWalkerHardcore-v2

1.2 Checkpoint 1

Start with the Pendulum-v0. Implement appropriate features for Linear function approximation Q-Learning. Make sure your feature evaluation is fast. Transform the continuous actions into a few discrete actions.

1.3 Checkpoint 2

Implement a different RL algorithm, Actor Critic or PGPE. Also consider implementing a experience replay to make use of previous runs. Test on Lunar Lander. (The

1.4 Final

Implement and analyze different methods on Lunar Lander and BipedalWalker-v2.