Assignment 13 (Bonus)

Machine Learning, Summer term 2013, Norman Hendrich

Solutions due by July 15

Exercise 13.1 (Random-Walk task, 3 points) Why do you think that a short (5 states) random-walk task was chosen for TD(0), but a longer one (19 states) for the demonstration of on-line and off-line TD(λ)?

Exercise 13.2 (Tabular case TD(λ), 3 points) Show that table-lookup TD(λ) is a special case of general TD(λ) (with function approximation).

Exercise 13.3 (Implement RL-learning with radial-basis functions, 14 points) Implement a RL learning tasks of your own choice that involves function-approximation with radial-basis functions. For example, another grid-world, a game, or a physics model like the pole-balancing cart. Document the chosen task, the selected state- and action-spaces, and the performance of your learner.