

KL learning for a Rat

Raoul-Martin Memmesheimer, Sep Thijssen, Bert Kappen

March 10, 2014

This is a computer exercises about KL-learning for a rat. Please write a short report in which you answer all the questions posed in the exercises below. Preferred format is Portable Document Format (.pdf) or plain text (.txt). Also please hand in source code files that you made or edited to perform the exercises, and make sure that these source files actually run without error and within a reasonable amount of time (in the order of 10 seconds or so).

To help you start up, we provide some Matlab code. The script `kl_definition.m` constructs a matrix \mathbf{q} , which defines the uncontrolled dynamics. It also constructs a vector \mathbf{r} that gives the state based cost. Together this pair defines a KL-control problem.

Exercise 1 (Uncontrolled search). Check numerically that the implementation of the transition matrix $q(x'|x)$, which can be found in `kl_definition.m`, yields a flat distribution for the number of visits at each point in space. Can you give an intuitive explanation why this is the case?

Exercise 2 (Optimal behavioral control). Compute the exact optimal KL-control transition matrix $p(x'|x)$ by solving the Perron-Frobenius eigenvalue problem for λ and β . Plot $-\log \beta$. Explain the interpretation of λ and β in the context of the control problem.

Exercise 3 (Test of the optimal behavioral control). Let your "rat" explore the environment according to the exact p and compare with the uncontrolled search.

Exercise 4 (Learning through experience). Simulate how a rat will determine $p(x'|x)$ if it does KL-learning and explores the environment according to $q(x'|x)$. Compare this learned approximate p with the exact p from exercise 2.

Exercise 5 (Maze). THIS IS A BONUS EXERCISE. Construct a maze by changing q and/or r and have your rat explore it.