

COMP 4106 - ARTIFICIAL INTELLIGENCE
WINTER 2018

ASSIGNMENT #3

DUE DATE: MARCH 29, 2018

Assignment Involving Elementary Reinforcement Learning

Introduction

The goal of this assignment is to have you implement some simple Learning Automata (LA) and Reinforcement Learning (RL) strategies. The application domain is quite straightforward, but it is typical of the domains where LA and RL can be used.

Problem Statement

Consider a satellite disk as in Figure 1 that receives data transmitted from a source. The satellite disk can rotate to 8 different location directions so as to maximize the strength of the received transmitted signal.




Figure 1: A simple model of a satellite disk which can be pointed in one of 8 possible directions.

The strength for each direction is unknown to the LA/RL scheme. However, it can be quantified as follows. Every index $i \in \{1, 2, \dots, 8\}$ is mapped to a unique number $G(i) \in \{1, 2, \dots, 8\}$ by a one-to-one and onto mapping. Then, the strength obeys the equation:

$$f_i = \frac{3G(i)}{2} + h, \quad i = \{1, 2, \dots, 8\},$$

where the signals are affected by a noise, h , that has a Gaussian distribution with a mean of 0 and a variance of an input parameter, σ^2 . In other words, at any time t , h is a random number generated¹ from a Gaussian distribution that represents the noise in the transmission.

¹You can use a Gaussian Random Number Generator from a standard existing library available in the language that you are using.

Assignment Objectives

These are the tasks you have to do:

- Learn the best direction for the disk using the Tsetlin, Krinsky, Krylov and L_{RI} schemes.
- In each case, use a suitable value for the “memory” and learning parameter.
- In each case, plot the ensemble average of the received signal strength for an ensemble of 100 experiments.

Questions

During the demo you should be prepared to discuss the following questions:

- Explain the way you chose the parameters for each scheme.
- Can you rank the schemes in terms of their speed/accuracy?