

Exercise sheet 9

Intelligent Agents

The following exercises are taken from “Artificial Intelligence: A Modern Approach”, Stuart J. Russell, Peter Norvig, 1995. <http://www.cs.berkeley.edu/~russell/aima1e/chapter02.pdf> provides some helpful information for the exercises.

Exercise 1 - Simulator (15 points)

Implement a performance-measuring environment simulator for the vacuum-cleaner world. This world can be described as follows:

- **Percepts:** Each vacuum-cleaner agent gets a three-element percept vector on each turn. The first element, a touch sensor, should be a 1 if the machine has bumped into something and a 0 otherwise. The second comes from a photosensor under the machine, which emits a 1 if there is dirt there and a 0 otherwise. The third comes from an infrared sensor, which emits a 1 when the agent is in its home location and a 0 otherwise.
- **Actions:** There are five actions available: go forward, turn right by 90°, turn left by 90°, suck up dirt, and turn off.
- **Goals:** The goal for each agent is to clean up and go home. To be precise, the performance measure will be 100 points for each piece of dirt vacuumed up, minus 1 point for each action taken, and minus 1000 points if it is not in the home location when it turns itself off.
- **Environment:** The environment consists of a grid of squares. Some squares contain obstacles (walls and furniture) and other squares are open space. Some of the open squares contain dirt. Each “go forward” action moves one square unless there is an obstacle in that square, in which case the agent stays where it is, but the touch sensor goes on. A “suck up dirt” action always cleans up the dirt. A “turn off” command ends the simulation.

Exercise 2 - Agent (15 points)

Implement a table-lookup agent for the special case of the vacuum-cleaner world consisting of a 2x2 grid of open squares, in which at most two squares will contain dirt. The agent starts in the upper left corner, facing to the right. Recall that a table-lookup agent consists of a table of actions indexed by a percept sequence. In this environment, the agent can always complete its task in nine or fewer actions (four moves, three turns, and two suck-ups), so the table only needs entries for percept sequences up to length nine. At each turn, there are eight possible percept vectors, so the table will be of size 8^9 . Fortunately, we can cut this down by realizing that the touch sensor and home sensor inputs are not needed; we can arrange so that the agent never bumps into a wall and knows when it has returned home. Then there are only two relevant percept vectors, ?0? and ?1?, and the size of the table is at most 2^9 . Run the environment simulator on the table-lookup agent in all possible worlds (how many are there?). Record its performance score for each world and its overall average score.