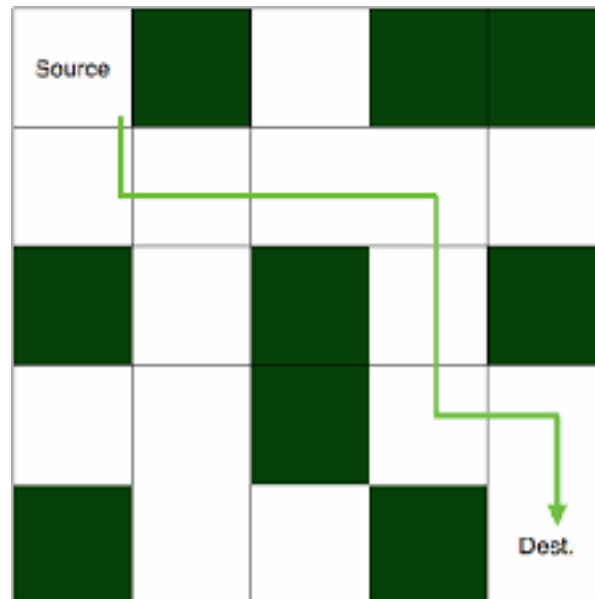


Artificial Intelligence

Practical Sheet 7: Reinforcement Learning



1. Create a Python script that generates random “N x N” boards, with each cell having only two possible types: “Safe” (penalty=1) or “Dangerous” (penalty=100).
2. Generate two random cell positions (each one provided in a (row, column) format), corresponding to the current position of the agent and to the final position.
3. Using Reinforcement Learning techniques (Q-learning), design and implement a solution for moving from the current state to the final position in as few movements as possible, avoiding the Dangerous cells as much as possible.
4. Compare different strategies for defining the “state” in terms of the spatial computational cost of the algorithm (the amount of storage required) and the quality of the solutions generated (the cost of the overall path, i.e., the sum of the penalties for the cells composing the best path).
 - a. Perceive the variations in the results, with respect to the α, γ values used in the Q-Table update formula:
5. Prepare a small report, describing your main findings.