15780: GRADUATE AI (SPRING 2018)

Homework 0: A^* Search

Release: January 17, 2018,
Due: January 26, 2018, 11:59pm

1   Programming Component [50 points]

We will use scientific Python for the implementation portions in this course. If you have not used Python before, we recommend downloading the Anaconda distribution (`https://www.continuum.io/downloads`) and looking through introductory resources like Google's Python Class (`https://developers.google.com/edu/python/`) and the Python Beginner's Guide (`https://wiki.python.org/moin/BeginnersGuide`). If you have not used scientific Python before, we recommend following introductions to NumPy (`http://www.numpy.org/`) and matplotlib (http://matplotlib.org/).

All Python code submissions in this course use Python 3.

For this problem, you will implement the A^* graph search algorithm to discover the series of moves that transform a moving tile puzzle from an initial state into a desired goal state. For example, for a 3×3 puzzle, given the following initial state:

$$\begin{array}{ccc} 1 & 2 & 3 \\ 4 & & 6 \\ 7 & 5 & 8 \end{array}$$

and the following goal:

$$\begin{array}{ccc} 1 & 2 & 3 \\ 4 & 5 & 6 \\ 7 & 8 & \end{array}$$

your program should return something like ["down", "right"], i.e., to solve this puzzle the blank has to be first moved down and then moved to the right.

We have provided a `search.py` Python module for you to get started with. When you are done, you will submit your completed module to Autolab for automatic grading. Do not rename the file or change the function names because our grader will import the module and functions by name. You are welcome to introduce new auxiliary functions and call them from the main functions we provide.

You should be able to execute the `search.py` we have provided without any modifications as a starting point. `search.py` contains a `main` function that shows expected outputs of every function we will grade.

1.1 Implementing Heuristics

The board state in Python will be a flattened tuple of the two-dimensional grid in row-major order where the blank tile is represented as 0. The example above is represented as (1, 2, 3, 4, 0, 6, 7, 5, 8). You can internally change the state to a format that's easier to work with if you prefer.

You should first implement the following two heuristics in the stubbed `heuristic_misplaced` and `heuristic_manhattan` functions we have provided.

1. **The number of misplaced tiles.** In the example above, there are two misplaced tiles, 5 and 8. The blank space is not tile and should not be included in your misplaced tile count.

2. **The sum of the Manhattan distances from the misplaced tiles to their correct positions.** In the example above, the distances from the misplaced tiles 5 and 8 to their correct positions are both 1, so the summed Manhattan distance is 2.

1.2 Implementing A^*

Next you will implement the A^* *graph search* algorithm in the stubbed `astar` function we have provided to find the shortest path using the heuristics above. Your function should return a string representing the moves needed to reach the goal and a list of states in the order they were expanded. Use the characters 'r', 'l', 'u', and 'd' for 'right', 'left', 'up', and 'down' directions, respectively. In the example above, your function should return the string 'dr'. If the grid is not valid (i.e., there is no valid path from the start state to the goal), return `None` for the optimal path, and the list of states expanded.

You should try your program on a number of puzzles with different initial states. Because the goal state cannot be achieved from all possible states generated by randomly placing the tiles on the board, you should write a function that shuffles a puzzle from the goal state to an initial state by repeatedly moving the blank to a position randomly chosen from the possible moves. The depth of the solution for your shuffled puzzle will be no greater than the number of times the blank is moved.

Because states can repeat, you may also want to maintain a separate list of all explored states, and only add nodes to the list if they have not already been explored. A^* is non-deterministic when selecting what node to explore next when they all have the same f value. To make A^* deterministic for grading, **we require that you break ties by selecting the lexicographically first node** based on the flattened state. For example, the flattened state (1, 2, 3, 4, 0, 6, 7, 5, 8) comes lexicographically before (1, 2, 3, 4, 0, 6, 7, 8, 5).

We recommend that you use the `heapdict` package as a priority queue to find elements with minimum cost. The `heapdict` package is included by default in Anaconda distributions and the project page is available at `https://github.com/DanielStutzbach/heapdict`. To break ties, your should use a tuple with (f value, state, other information you want to store) as the value in your heapdict.

2 Written Component [50 points]

Heuristics for n-puzzle

In class (lecture 2), we examined the 8-puzzle game and two heuristics. In this question, we generalize the puzzle and consider a different heuristic.

The 8-puzzle game is a single-player board game which consists of a 3×3 board with 8 tiles and 1 blank slot. A tile can move horizontally or vertically to its adjacent blank slot (if it neighbors it). The objective of the game is to start from a given arrangement of tiles and move tiles to achieve a given goal arrangement. Now, consider the n-puzzle game which consists of an $n \times n$ board with $n^2 - 1$ tiles and 1 blank slot. The rules for the n-puzzle game are exactly the same as those for the 8-puzzle. The following questions are about the n-puzzle game.

Here, we discuss some heuristics to be used with A^* graph search. Recall that heuristic $h_1(\cdot)$ returns the number of tiles that are in the wrong position and $h_2(\cdot)$ returns the sum of Manhattan distances of tiles from their goal positions. We introduce a third heuristic, $h_3(\cdot)$, that is the *minimum* number of moves necessary to get to the goal state if each action could move any tile to the blank slot. This is another *relaxed problem* heuristic.

1. (15 points) Prove that h_3 is consistent.

2. (5 points) Prove or disprove: h_3 dominates h_1.

3. (5 points) Prove or disprove: h_3 dominates h_2.

4. (5 points) Prove or disprove: the heuristic $h = \max(h_2, h_3)$ is consistent.

5. (20 points) An important feature of any heuristic is that it can be computed efficiently. Give a polynomial-time algorithm in the number of tiles to compute h_3 for a given state. Prove its correctness and running time.

3 Submitting to Autolab

Create a tar file containing your writeup for the second problem and the completed `search.py` module for the first problem. Make sure that your tar has these files at the root and not in a subdirectory. Use the following commands from a directory with your files to create a `handin.tgz` file for submission.

```
$ ls
search.py  writeup.pdf
$ tar cvzf handin.tgz writeup.pdf search.py
a writeup.pdf
a search.py
$ ls
handin.tgz  search.py  writeup.pdf
```