


1 Tower of Hanoi



The Tower of Hanoi is a canonical puzzle studying problem solving and formulation. The puzzle starts with n disks of different sizes stacked in order of size (see picture above) on a peg, along with two empty pegs. We can move disks freely between the pegs, but larger disks cannot be stacked on top of smaller ones. The goal is to move all disks to the third peg.

We will attempt to formulate Tower of Hanoi as a search problem.

(a) How could we represent this puzzle as a problem, ie. what would the states be?

Possible answer: store 3 lists tracking the disks stacked on each peg. Enumerate the disks 1 to n .

For the following questions, assume that you have used your state representation in your answer above.

(b) What is the size of the state space in terms of n ?

Assume that a disk number can appear on any of the three lists, independently of where the other numbers are. This means that there are 3^n possible combinations of the three lists. We don't need to factor in the order of those lists, because the only legal order is sorted from smallest to largest. For example, with only two disks, the number of combinations is 3^2 . Specifically:

- ([1, 2], [], [])
- ([1], [2], [])
- ([1], [], [2])
- ([2], [1], [])
- ([], [1, 2], [])
- ([], [1], [2])
- ([2], [], [1])
- ([], [2], [1])
- ([], [], [1, 2])

(c) What is the starting state?

([1, 2, ..., n], [], [])

(d) From some given state, what legal actions are there?

We can pop the first integer from any list and push it to the front of any other list, given that this integer is less than the current first integer of the target list (or that the target list is empty).

(e) What is the goal test? Remember that this determines whether a given state is a goal state, `goalTest(state)`. `goalTest` should check that `state == ([, [], [1, 2, ..., n])`.

2 Nim Continued

Today we get to bring the Nim agents from lecture to life! Download and unzip from the Recitation page on the course website:

- <https://www.cs.cmu.edu/~15281/#recitations>

This code is meant for you to mess around with during recitation. Don't worry about breaking things, as you can always redownload the original code. Inside the Nim directory you'll find:

- `nim.py`: The game environment
- `agent.py`: The Agent base class. All agents inherit from this simple base class.
- `agent*.py`: Implementations of various agents. We saw some of these in lecture.
- `play_one.py`: Simple script to have two agents play one round of the game
- `play_many.py`: Simple script to have two agents play many rounds of the game

Feel free to edit any of these files or make copies of `play*.py` or `agent*.py` to create your own battles or agents. But first, we'll have you just run and inspect the code to get an idea of how things are working in this simple game. You can start by running the following:

- `python play_one.py`
- `python play_many.py`

Take a look through the source code and answer the following questions:

(a) How is the state represented in the code?

[The number of remaining pieces](#)

(b) What is Agent005's strategy?

[Choose a random action](#)

(c) What is Agent003's strategy?

[Copy the opponent's action](#)

(d) How do the AgentRL stats change when playing against an Agent008 rather than an Agent007? What difference in the two agents do you think causes this? (You should modify `play_many.py`)

[With an Agent008, fewer of the table entries have N/A's. This is because Agent008 will act randomly when neither of its actions is optimal \(it can't force its opponent to have mod 3 pieces\). Because of this randomness, the RL agent will see more game states than it might have versus an Agent007.](#)

(e) When we run the games 1000 and 10000 times, how do the AgentRL stats change when playing against another AgentRL? Which states seem to be good states to take? (You should modify `play_many.py`)

[The stats do not change that much since the two agents are both RL agents under explore mode \(random actions\), and neither agent is exploiting their table of statistics to try to win more games. For example, looking at AgentRL table of statistics, action 2 from state 5 has a much higher win probability \(75%\) than action 1 from state 5 \(35%\). However, because explore mode defaults to true, AgentRL will continue choosing random actions from state 5, rather than the more attractive action 2.](#)

(Bonus) AgentRL has an `explore_mode` setting that defaults to `True`. What happens when you play 500 games, turn off explore mode (`bb8.explore_mode = False`), and then play 500 more games?

The `explore_mode` indicates whether or not the agent should explore - act randomly, so that it can record as many different game states into its internal table as possible, or exploit - where it uses the information it has learned to play as well as it can. Since `explore_mode` is initially true, the RL agent is only acting randomly. However, once it uses its table, it begins to do much better, and win a lot more of its games. This is a concept we'll come back to when we learn about reinforcement learning (RL) :)