

— GUIDE —

Assignment 3: Wumpus Quest

AI-2 Systems Project (Summer Semester 2026)

Jan Frederik Schaefer

Friedrich-Alexander-Universität Erlangen-Nürnberg, Department Informatik

This document is intended to help you solve the assignment “Assignment 3: Wumpus Quest” [AS]. You do not have to read it, but we do recommend to at least take a look at the tips and common issues.

1 A few tips

1. Consider using value iteration.
2. Model the rewards in a way that the utilities roughly correspond to the expected amount of gold that you will carry out of the cave. This way, maximizing the utility corresponds to maximizing the expected amount of gold, i.e. maximizing the rating of your agent. You might have to deviate from this a bit, but it should be a good starting point.
3. It can help with debugging if you can somehow look at the computed utilities/policy.
4. Use Markov Decision Processes for the first environment (even though it can be solved without) to get some experience before moving on to the harder environments.
5. Code optimizations:
 - If you cache the utilities/policy, you do not have to re-compute them for every action.
 - The number of possible states is relatively large – try to avoid making it much larger than necessary.
 - Value iteration (or similar algorithms) might be much faster if you once create an efficient representation of the state space and the transition model, and then only use that representation. For example, you could assign an integer to each state.
 - As a typical run involves many actions and the delays from the interaction with the server may become problematic. To mitigate this, you can have several runs in parallel. The server then sends you multiple action requests within a single

HTTP request (and expects that many responses). If you use the example client, you can easily switch to that behaviour.

6. The server interaction is a bit tricky. Please reach out if you face any difficulties or have ideas how it could be improved.

2 Common issues

- The agent doesn't want to leave the cave. It can help to have a small negative reward for each step in the cave to nudge the agent towards the exit.
- Not using MDPs. You can quickly get first results without using MDPs. However, past experience shows that non-MDP approaches usually cannot get particularly good scores. The reason might be that the assignment was set up with MDPs in mind.
- Not setting up the rewards in a way that utilities correspond to gold. You can still get many points without this, but most people struggled to get full points without doing this.

References

- [AS] *Assignment 3: Wumpus Quest*. URL: <https://kwarc.info/teaching/AISysProj/SS26/assignment-2.3.pdf>.