

Assignment 0 (Warm-Up, Variant A): Find Back to the Wumpus Cave

AI-2 Systems Project (Summer Semester 2026)

Jan Frederik Schaefer

Friedrich-Alexander-Universität Erlangen-Nürnberg, Department Informatik

Topic: Basic probabilities
Deadline: May 16, 2026
Extended deadline: July 18, 2026
Version from: April 28, 2026
Author: Jan Frederik Schaefer

Important notes:

- Ask for help if you are stuck (office hours, assignment room, ...).
- Every assignment has a guide with tips – you can find it at [\[AG\]](#).
- The extended deadline is granted automatically.
- The extended deadline is after the deadlines of some other assignments, i.e. submitting at the extended deadline means you cannot finish them this semester.

**This assignment has to be solved individually (not as a team).
Using someone else's solution code, even as inspiration, is not allowed!
Sharing your solution code with other AISysProj students is not allowed.**

1 Task summary

The Wumpus, a mythical creature that repeatedly shows up in the AI systems project, has left its cave and got lost. Your task is to make a travel plan for the Wumpus. Ideally, this travel plan will get the Wumpus back to the cave soon, but it is not a requirement – any travel plan will be accepted.

Your implementation will get different problem instances from the AISysProj server and has to send back its travel plan. That lets the server evaluate your agent and can help you with debugging. The travel plan will be evaluated based on the how likely the travel plan is to succeed and how long it would take to reach the cave (details in Section [2.3](#)). The

assignment repository [AR] has a script that already implements the server interaction for you.

Didactic objectives

1. Develop an algorithm to solve a non-trivial problem,
2. implement a small software project from scratch,
3. get hands-on experience working with basic concepts from probability theory,
4. get to know the AISysProj setup and workflows.

Prerequisites and useful methods

1. The basics of probability theory: condition probabilities, independence, expected values, Bayes' rule, ...

2 Navigating the Wumpus world

The Wumpus got lost. Your task is to implement an AI agent that makes a plan for it to get back to the Wumpus cave. We have a server to help evaluate your agent. It will send you a map along with some other information and your agent will have to respond with a plan to find back to the cave.¹ The server has different environments, corresponding to different difficulty levels. For the easier environments, not everything mentioned in this section is relevant.

2.1 Maps

Maps of the Wumpus world are encoded as a string, where each line corresponds to a row of cells and each character describes the properties of an individual cell. Here is the string that describes the map in Figure 1:

```
SBCMW  
WWCMM  
BBCBM  
WBCMM  
MBCSM
```

¹The server part may seem daunting, but we provide an example implementation in Python. All you have to then do is implement a function.

S <i>P</i> : 0.0000 <i>T</i> : -	B <i>P</i> : 0.0203 <i>T</i> : -	C <i>P</i> : 0.0406 <i>T</i> : -	M <i>P</i> : 0.0000 <i>T</i> : -	W <i>P</i> : 0.0000 <i>T</i> : -
W <i>P</i> : 0.0000 <i>T</i> : -	W <i>P</i> : 0.0000 <i>T</i> : -	C <i>P</i> : 0.0406 <i>T</i> : -	M <i>P</i> : 0.0000 <i>T</i> : -	M <i>P</i> : 0.0000 <i>T</i> : -
B <i>P</i> : 0.1624 <i>T</i> : 0.5000	B <i>P</i> : 0.1624 <i>T</i> : 0.5000	C <i>P</i> : 0.0406 <i>T</i> : 1.5000	B <i>P</i> : 0.1624 <i>T</i> : -	M <i>P</i> : 0.0000 <i>T</i> : -
W <i>P</i> : 0.0000 <i>T</i> : -	B <i>P</i> : 0.1624 <i>T</i> : 2.5000	C <i>P</i> : 0.0406 <i>T</i> : 2.5000	M <i>P</i> : 0.0000 <i>T</i> : -	M <i>P</i> : 0.0000 <i>T</i> : -
M <i>P</i> : 0.0000 <i>T</i> : -	B <i>P</i> : 0.1624 <i>T</i> : 1.5000	C <i>P</i> : 0.0051 <i>T</i> : 3.5000	S <i>P</i> : 0.0000 <i>T</i> : -	M <i>P</i> : 0.0000 <i>T</i> : -

Observations: {"humidity": 0, "current-cell": "B"}

Actions: GO north, GO west, GO north, GO north, GO north, GO east

Success chance: 73.60406091%

Expected duration: 1.34828

Figure 1: An example map of the Wumpus world of one of the more difficult environments. You can find similar visualizations for all environments on the AISysProj server. The P values are the probabilities that the Wumpus is in a particular cell given the example observations. The T values indicate the time it takes to reach the cave from that cell (“-” if the cave is not reachable) when using the example plan.

We have the following cell types:

- M: Meadows.
- B: Trees (broad-leaf).
- C: Trees (coniferous).
- S: Swamps.
- W: Wumpus cave entrance. Your goal is to reach one of those.

We assume that everything outside the map are meadows (M).

2.2 Observations

We do not know where the Wumpus is (after all, the Wumpus got lost). Therefore, we assume that, a priori, the Wumpus is equally likely in any of the cells on the map. However, the Wumpus can make some basic observations to narrow down the possibilities:

- It observes what type of cell it currently is in (M or B or ...). Unfortunately, the Wumpus has bad vision and misidentifies trees 20% of the time (i.e. if the Wumpus is in a B cell, there is a 20% chance that it instead thinks it is a C cell and vice versa).
- In the more advanced environments, the Wumpus also perceives the humidity. The base humidity score is 0. Each swamp in the four cells adjacent to the Wumpus increases the humidity score by 1. The score is additionally increased by 2 if the Wumpus is currently in a swamp. As the Wumpus cannot perceive the humidity accurately, there is a 10% chance that the humidity score is 1 too low and a 10% chance that it is 1 too high. The probability that the reported score is accurate is therefore 80%.

2.3 Plans and their execution time

Your task is to make a plan for the Wumpus. A plan is a sequence of actions of the form `GO [north|east|south|west]`. For each action, the Wumpus goes one cell in the indicated direction. Note that the world is not limited to the map and the Wumpus can go to cells outside the map (recall that every cell outside the map is a meadow).

Ideally, the Wumpus reaches one of the Wumpus cave entrances while following the plan. How long it takes for the Wumpus to reach the entrance depends on the cell types. We assume that the Wumpus starts in the center of a cell and only has to reach the edge of a cell with a cave entrance. Usually, crossing a complete cell takes 1 hour, but crossing a swamp cell takes 2 hours. Additionally, the Wumpus has to put on Wellington boots when entering a swamp and take them off again when leaving a swamp, each of which takes

an additional hour (the Wumpus has many feet). Note that the Wumpus does not take the boots off and put them on again when travelling from one swamp cell to the next. Furthermore, the Wumpus is already wearing boots if it starts in the swamp and does not take them off when reaching the cave. Let us consider the following map as an example:

```
CS
```

```
MW
```

If the Wumpus is initially the **C** cell and follows the plan **GO east**, **GO south**, then it will have to cross $\frac{1}{2}$ of the **C** cell (as it starts in the center), which takes 0.5 hours, and $\frac{2}{2}$ of the **S** cell ($\frac{1}{2}$ to reach its center and $\frac{1}{2}$ to reach the edge of the **W** cell), which takes 2 hours. Additionally, the Wumpus will have to put on boots when entering the **S** cell, which takes an additional hour. In total, it will therefore take $0.5 + 1 + 2 = 3.5$ hours to reach the cave.

See Figure 1 for a more complex example.

2.4 Success chance and expected time

Along with the plan, your agent should also provide the probability of the plan succeeding and the expected time it will take to reach the cave. Let us consider the following map as an example:

```
BWM
```

```
MBB
```

```
BMW
```

Let us assume that we know that the Wumpus is in a **B** cell. Every environment has a maximum time for the Wumpus to reach the cave. Let us assume that it is 3 hours for this example.

Now, the agent has to create some plan. Any plan is acceptable, but plans with a higher success chance and a lower expected time are better. We will use: **GO north**, **GO east**, **GO south**, **GO east**.

If the Wumpus is in the top left **B** cell, it will take 2.5 hours to reach a **W** cell (recall that the Wumpus can go outside the map). If the Wumpus is in the center **B** cell, it will take 0.5 hour to reach a **W** cell. If the Wumpus is in the right **B** cell, it will never reach a **W** cell. If the Wumpus is in the bottom left **B** cell, it would theoretically take 3.5 hours to reach a **W** cell, but as the maximum time is 3 hours, the plan would be considered unsuccessful in this case.

Now, the **success chance** is the probability that the Wumpus is in a cell, for which the plan is successful. In this case, the Wumpus is equally likely in each **B** cell, and the plan

succeeds for two of them, so the success chance is $\frac{2}{4}$.

The **expected time** is the expected value of the time it takes to reach the cave, only considering the cases where the plan is successful. In this case, it would be $\frac{1}{2} \cdot 2.5 + \frac{1}{2} \cdot 0.5 = 1.5$ hours.

3 Evaluation on the server

You should evaluate your implementation with the AISysProj server: <https://aisysproj.kwarc.info/>. You will get JSON requests from the server and have to respond with your plan. For example, a request could have the following content (in this assignment you can ignore the `initial-equipment`):

```
{
  "map": "SBCMW\nWWCMM\nBBCBM\nWBCMM\nMBCSM",
  "observations": {"humidity": 0, "current-cell": "B"},
  "initial-equipment": [],
  "max-time": 5.0
}
```

Your agent should then respond with a plan, the probability of the plan succeeding, and the expected time it will take to reach the cave:

```
{
  "actions": [
    "GO_west",
    "GO_north",
    "GO_south"
  ],
  "expected-time": 1.1428571428571428,
  "success-chance": 0.5685279187817259
}
```

The server will then evaluate your response and give you a rating. If the success chance and the expected time are correct, your rating will be computed by the formula

$$\text{rating} = \text{success-chance} \cdot \text{expected-time} + (1 - \text{success-chance}) \cdot \text{max-time}$$

where `max-time` is the maximum time for the environment.

The server remembers your last 1000 responses and computes the average rating (the lower, the better). If you computed the wrong value for `success-chance` or `expected-time`, the server will instead use `max-time` for the rating. However, afterwards you can see the problem on the server website with the correct expected time and some other information that could help you with debugging.

The server has different environments that you can use (some are easier than others). Your repository for this assignment will contain one configuration file for each environment, which contains, among other things, your agent name and a password. The assignment repository [AR] contains an example implementation in Python along with instructions on how to use it, so that you do not have to worry about the technical aspects of communicating with the server. If you want to use a different programming language, you can find the protocol specification at [CSP]. We are happy to help you with implementing it if you are prepared to donate your code for others.

4 What to submit

You should push your solution to your git repository for this assignment. Concretely, your repository should contain:

1. all your code for solving this assignment,
2. a `README.md` file explaining
 - i. dependencies (programming language, version, external libraries and how to get them),
 - ii. how to run your code,
 - iii. the repository structure,
 - iv. anything else we should know,
3. a solution summary (see [SoS] for more details – it should describe the main ideas, not document the code),
4. an empty file `submitted.txt` to indicate that you have submitted your solution.

Furthermore, you should run your code so that the server has an evaluation for your agent.

5 Points and environments

The total number of points for this assignment is 100. You can get up to 80 points for the performance of your agent according to the server. Each environment allows you to get a

certain numbers of points if your agent performs well enough. When grading, we will only consider the environment in which you would get most points (we do not add up results from different environments). That means that you do not have to run your code on every environment and you can get full points if you only run it on the most difficult one where you can get up to 80 points.

Additionally, you can get up to 20 points for the quality of the submission (README, evaluation, ...). You can only get points for the submission quality if you get at least 1 point for the performance of your agent.

Note: It is okay if your agent is a bit “lucky” and gets a slightly higher rating than it usually does. We will use the best rating on the server for your grade, assuming that we can reproduce a similar performance (i.e. if you get a 5.39 rating on the server and we get a 5.48 rating that is okay – but if you get a 5.39 rating on the server and we get a 6.79 rating when testing it, we might ask some questions).

Below is a table that summarizes the properties of the different environments and how many points you can get for each. For each environment, you have a config file.

Config file	Map size	Cells	Max. time	Equipment	Observations	Points
env-1.json	5 × 5	C M	2	–	current-cell	$\left\{ \begin{array}{l} 0 \quad \text{if rating} > 1.8 \\ 20.0 \quad \text{if rating} \leq 1.8 \\ 30.0 \quad \text{if rating} \leq 1.5 \end{array} \right.$
env-2.json	5 × 5	C M S	2	–	current-cell	$\left\{ \begin{array}{l} 0 \quad \text{if rating} > 1.8 \\ 30.0 \quad \text{if rating} \leq 1.8 \\ 40.0 \quad \text{if rating} \leq 1.7 \end{array} \right.$
env-3.json	5 × 5	C M S	5	–	current-cell	$\left\{ \begin{array}{l} 0 \quad \text{if rating} > 4.5 \\ 40.0 \quad \text{if rating} \leq 4.5 \\ 50.0 \quad \text{if rating} \leq 3.5 \end{array} \right.$
env-4.json	5 × 5	B C M S	5	–	current-cell	$\left\{ \begin{array}{l} 0 \quad \text{if rating} > 4.5 \\ 45.0 \quad \text{if rating} \leq 4.5 \\ 55.0 \quad \text{if rating} \leq 3.5 \end{array} \right.$
env-5.json	5 × 5	B C M S	5	–	current-cell, humidity	$\left\{ \begin{array}{l} 0 \quad \text{if rating} > 4.5 \\ 50.0 \quad \text{if rating} \leq 4.5 \\ 60.0 \quad \text{if rating} \leq 3.5 \\ 70.0 \quad \text{if rating} \leq 3.0 \end{array} \right.$
env-6.json	15 × 15	B C M S	24	–	current-cell, humidity	$\left\{ \begin{array}{l} 0 \quad \text{if rating} > 20.0 \\ 50.0 \quad \text{if rating} \leq 20.0 \\ 65.0 \quad \text{if rating} \leq 18.0 \\ 75.0 \quad \text{if rating} \leq 16.0 \\ 80.0 \quad \text{if rating} \leq 14.5 \end{array} \right.$

References

- [AG] *Guide for “Assignment 0 (Warm-Up, Variant A): Find Back to the Wumpus Cave”*. URL: <https://kwarc.info/teaching/AISysProj/SS26/assignment-2.0-guide.pdf>.
- [AR] *Repository for Assignment 0 (Warm-Up, Variant A): Find Back to the Wumpus Cave*. URL: <https://gitlab.rrze.fau.de/wrv/AISysProj/ss26/a2.0-find-wumpus-cave/assignment>.
- [CSP] Jan Frederik Schaefer. *AISysProj server – Clients and server protocol*. URL: <https://aisysprojserver.readthedocs.io/en/latest/clients.html>.
- [SoS] *Solution Summary*. URL: <https://gitlab.rrze.fau.de/wrv/AISysProj/admin/general/-/blob/main/solution-summary.md>.