Einführung in Artificial Intelligence SS 2025, 4.0 VU, 192.027

Exercise Sheet 1 – Agents and Search I

For the discussion part of this exercise, mark and upload your solved exercises in **TUWEL** until Wednesday, June 11, 23:55 CEST. The registration for a solution discussion ends on Friday, June 13, 23:55 CEST. Be sure that you tick only those exercises that you can solve and explain!

In the discussion, students will be asked questions about their solutions of examples they checked. The discussion will be evaluated with 0–25 points, which are weighted with the fraction of checked examples and rounded to the next integer. There is *no minimum number of points* needed for a positive grade (i.e., you do not need to participate for a positive grade, but you can get at most ≈80% without doing exercises).
Note, however, that *your registration is binding*. Thus, *if* you register for a solution discussion, then it is *mandatory* to show up. Not coming to the discussion after registration will lead to a reduction of examination attempts from 4 to 2.
Please ask questions in the **TUWEL** forum or visit our tutors during the tutor hours (see **TUWEL**).

Exercise 1.1: Consider the following graph:

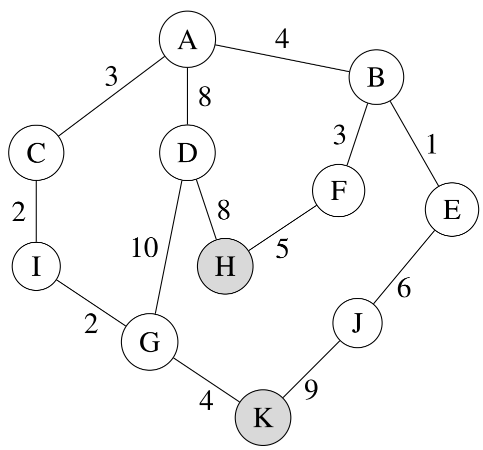

Let A be the initial state and the grey vertices the goals. Determine for the following search strategies the order in which the nodes are expanded and the corresponding goal node. In addition, compute for each search strategy the set of nodes that is actually kept in memory when the goal state is found (A has depth zero). When multiple nodes can be added to the frontier, add them in such a way that they are expanded in alphabetical order. When there is no risk of ambiguity, you can denote a node of the search tree by the state it has.

- Breadth First Search
- Uniform-Cost Search
- Depth First Search
- Depth-Limited Search (use a limit of 2)
- Iterative Deepening Search

Exercise 1.2: Look again at the graph of Exercise 1.1. Assume A is the initial and H the goal state (i.e. K is no more a goal). What problem do you encounter when applying *depth first search* (tree-like search version, as illustrated on the slides)? Discuss how the algorithm can be changed such that the problem is solved, and explain why your modification has no impact on the properties of the DFS which were discussed in the lecture.

Exercise 1.3: Decide and explain which of the following statements are true and which are false? Back up your answers with proofs or counterexamples.

(1) A task environment in which every agent is rational does not exist.

(2) There is a state space in which IDS performs significantly worse than DFS.

Exercise 1.4: In this exercise, we will see that some agents which have rich enough capabilities of *self-consciousness* cannot exist in principle. To this end, we define the notion of a *Gödelian agent*[1] as follows. Imagine such an agent to be a device which is able to tell us statements of a specific form. The statements the agent can tell us are build up using the following symbols:

$$\neg, T, N, (,)$$

We call the set of all statements which we can form using these symbols the *language* of our agent. For example, the statement $\neg T(T)$ is in the agent's language. The *norm* of a statement X is the statement $X(X)$. (Here, X denotes an arbitrary statement.) Not all statements in this language are meaningful, only sentences are. A statement is a *sentence* if it is of one of the following forms:

1. $T(X)$,
2. $\neg T(X)$,
3. $TN(X)$,
4. $\neg TN(X)$.

We now assign *truth values* to sentences as follows:

1. $T(X)$ is true iff X can be told by the agent;
2. $\neg T(X)$ is true iff X cannot be told by the agent;
3. $TN(X)$ is true iff the norm of X can be told by the agent;
4. $\neg TN(X)$ is true iff the norm of X cannot be told by the agent.

We assume our agent to be trustworthy, i.e., we assume that whenever the agent tells us a sentence, then it is true. Now your task is to show that, under this assumption, the opposite does not hold, i.e., prove that there is a true statement which cannot be told by our trustworthy agent. *Hint:* find a statement which is true iff *the statement itself* cannot be told by the agent. Use then the assumption of trustworthiness to conclude that your statement cannot be told. Think about why your statement cannot be told and discuss the reason(s). *Warning:* any "statement" involving X is wrong, for X is a metavariable, not really a symbol in the language of our agent.

[1] The self-referential character of this agent is very similar to the self-referential techniques used on the proofs of Gödel's seminal incompleteness results. Do not get scared, to solve this exercise you do not need to know anything about these.

Exercise 1.5: Let $b > 1$ be the maximal branching degree in the search tree and let d be its depth. Estimate the number of nodes, $n_{dfid}(d)$, generated during a depth-first-iterated-deepening search with depth d. Show that $n_{dfid}(d)$ is $O(b^d)$ and estimate the constant c_{dfid}. What can you say about the overhead induced by dfid?

Exercise 1.6: Give the PEAS description and the characteristics of the task environment for each of the following intelligent agents. Be sure to explain your reasoning and assumptions.

- Robot vacuum cleaner,
- Playing MMORPG such as *World of Warcraft*,
- An AI for the game Go.