

Foundations of Artificial Intelligence

T. Keller, F. Pommerening
S. Sievers
Spring Term 2023

University of Basel
Computer Science


Exercise Sheet 12

Due: May 28, 2023

Important: for submission, consult the rules at the end of the exercise. Non-adherence to the rules will lead to your submission not being corrected.

Exercise 12.1 (2+2 marks)

Consider the game inducing the following game tree T . Note that we excluded moves (the edges are unlabeled) because they are irrelevant for the purpose of this exercise.



- Apply alpha-beta search to T . Successor nodes should be considered from left to right. As a solution for this part, draw the pruned game tree (i.e., only the nodes that are considered by the algorithm) and annotate all nodes with the utility as well as the alpha and beta values (as they are when the algorithm terminates). How many nodes does alpha-beta search need to consider and how does this compare to minimax?
- Reorder T (without altering the game) such that alpha-beta search prunes the maximum possible amount of nodes. Draw the reordered game tree and again annotate all nodes with the utility and the alpha and beta values and show which parts of the tree can be pruned (e.g., by drawing cut lines through edges which lead to subtrees that do not need to be considered). How many nodes does alpha-beta search need to consider? How does this compare to alpha-beta search with the original order?

Exercise 12.2 (2 marks)

Consider the stochastic game inducing the following game tree T .

Exercise 12.4 (1+1 marks)

- (a) Consider a search node of a game tree with three successors n_1 , n_2 and n_3 and $e^{\frac{n_1 \cdot \hat{v}}{\tau}} = 60$, $e^{\frac{n_2 \cdot \hat{v}}{\tau}} = 30$ and $e^{\frac{n_3 \cdot \hat{v}}{\tau}} = 90$. Provide the probabilities that n_1 , n_2 and n_3 are selected by MCTS with softmax tree policy.

Hint: You do not need a calculator to answer this question.

- (b) Consider a search node of a game tree with two successors n_1 and n_2 and $n_1 \cdot \hat{v} = 10$, $n_1 \cdot N = 5$, $n_2 \cdot \hat{v} = 6$ and $n_2 \cdot N = 8$. Assume successor n_2 was selected in the last trial by MCTS with UCB1 tree policy. How have the state value estimates been in comparison to the current state value estimates? What does this mean for the reward of the last trial?

Submission rules:

- Exercise sheets must be submitted in groups of two students. Please submit a single copy of the exercises per group (only one member of the group does the submission).
- Create a single PDF file (ending .pdf) for all non-programming exercises. Use a file name that does not contain any spaces or special characters other than the underscore “_”. If you want to submit handwritten solutions, include their scans in the single PDF. Make sure it is in a reasonable resolution so that it is readable, but ensure at the same time that the PDF size is not astronomically large. Put the names of all group members on top of the first page. Either use page numbers on all pages or put your names on each page. Make sure your PDF has size A4 (fits the page size if printed on A4).
- For programming exercises, only create those code textfiles required by the exercise. Put your names in a comment on top of each file. Make sure your code compiles and test it. Code that does not compile or which we cannot successfully execute will not be graded.
- For the submission: if the exercise sheet does not include programming exercises, simply upload the single PDF. If the exercise sheet includes programming exercises, upload a ZIP file (ending .zip, .tar.gz or .tgz; *not* .rar or anything else) containing the single PDF and the code textfile(s) and nothing else. Do not use directories within the ZIP, i.e., zip the files directly.
- Do not upload several versions to ADAM, i.e., if you need to resubmit, use the same file name again so that the previous submission is overwritten.