



1 Adversarial Search

Today, we will be exploring different adversarial searches, such as Minimax and Alpha Beta pruning.



(a) Consider the following zero-sum game tree: trapezoids pointing up represent maximizers and trapezoids pointing down represent minimizers. It is currently your turn and both you and your opponent act optimally. Carry out the Minimax search algorithm and write the value of each node inside the trapezoids.



(b) What move should you make now? How much is the game worth to you?

The game is worth 1. We should make the move that takes us left down to the node containing 1.

(c) Now, consider the same game tree. Using alpha beta pruning and expanding successors from left to right, record the values of alpha and beta at each node. Furthermore, write the value being returned at each node inside the trapezoid. Put an 'X' through the edges that are pruned off.






(d) Which leaf nodes are never visited due to pruning?




6, 12, -3, -2, 10, -5, -7

2 Alpha Beta Expinimax

In this question, player A is a minimizer, player B is a maximizer, and C represents a chance node. All children of a chance node are equally likely. Consider a game tree with players A, B, and C. In lecture, we considered how to prune a minimax game tree - in this question, you will consider how to prune an expinimax game tree (like a minimax game tree but with chance nodes). Assume that the children of a node are visited left to right.

For each of the following game trees, give an assignment of terminal values to the leaf nodes such that the bolded node can be pruned (it doesn't matter if you prune more nodes), or write "not possible" if no such assignment exists. You may give an assignment where an ancestor of the bolded node is pruned (since then the bolded node will never be visited). You should not prune on equality, and your terminal values must be finite (including negative values).





Not possible. At the bolded node, the minimizer can guarantee a score of at most the value of its left subtree. However, since we have not visited all the children of C, there is no bound on the value that C could attain. So, we need to continue exploring nodes until we can put a bound on C's value, which means that we have to explore the bolded node.

3 True/ False Section

(a) Consider an adversarial game tree where the root node is a maximizer, and the minimax value of the game is V_M . Now, also consider an otherwise identical tree where every minimizer node is replaced with a chance node (with an arbitrary but known probability distribution). The expectimax value of the modified game tree is V_E .

True False V_M is guaranteed to be less than or equal to V_E .

True: The maximizer is guaranteed to be able to achieve v_M if the minimizer acts optimally. He can potentially do better if the minimizer acts suboptimally (e.g. by acting randomly). The expectation at changes nodes can be no less than the minimum of the nodes' successors.

True False Using the optimal minimax policy in the game corresponding to the modified (chance) game tree is guaranteed to result in a payoff of at least V_M .

True: The minimax policy is always guaranteed to achieve payoff of at least v_M , no matter what the minimizer does.

True False Using the optimal minimax policy in the game corresponding to the modified (chance) game tree is guaranteed to result in a payoff of at least V_E .

False: In order to achieve v_E in the modified (chance) game, the maximizer may need to change his or her strategy. The minimax strategy may avoid actions where the minimum value is less than the expectation. Moreover, even if the maximizer followed the expectimax strategy, he or she is only guaranteed v_E *in expectation*.

(b) Consider a one-person game, where the one player's actions have non-deterministic outcomes. The player gets +1 utility for winning and -1 for losing. Mark all of the approaches that can be used to model and solve this game.

- Minimax with terminal values equal to +1 for wins and -1 for losses
- Expectimax with terminal values equal to +1 for wins and -1 for losses
- Value iteration with all rewards set to 0, except wins and losses, which are set to +1 and -1
- None of the above

- Minimax with terminal values equal to +1 for wins and -1 for losses
- Expectimax with terminal values equal to +1 for wins and -1 for losses
- Value iteration with all rewards set to 0, except wins and losses, which are set to +1 and -1
- None of the above

Minimax is not a good fit, because there is no minimizer - there is only a maximizer (the player) and chance nodes (the non-deterministic actions). The existence of a maximizer and chance nodes means that this is particularly suited to expectimax and value iteration.