

1. [3 marks] For each person or program, give the label of its description.

Aja Huang _____ Cho Chikun _____ David Silver _____ Demis Hassabis _____

Fan Hui _____ Geoff Hinton _____ Lee Sedol _____ Michael Redmond _____

(a) former European go champion who lost to a strong computer program in a 5-game match in 2015 (b) CEO of a leading AI research company (c) lead programmer of AlphaGo (d) former #1 go player who lost to a strong computer program in a 5-game match in 2016 (e) 9dan professional go player and commentator (f) neural net expert who wrote a paper on image classification (g) first author of Nature paper on AlphaGo (h) former #1 go player who defeated a strong computer program in a 3-game match in 2016 .

2. [4 marks] Fill in the blanks, and circle correct answers.

AlphaGo integrates neural net calls into _____ search , overcoming the slowness of net calls by **(circle all that apply)** a) building shallow nets that reply quicker than the initial deeper versions b) handling net calls with GPUs c) having the master algorithm continue to operate while net calls are executing d) having each net call distributed over parallel processes .

AlphaGo child-selection uses **(circle all that apply)** a) a deep policy net b) a deep value net c) a shallow policy net d) a shallow value net e) simulations .

In AlphaGo, once a leaf is reached, using a **(circle all that apply)** cpu / gpu , a call is made on a **(circle all that apply)** a) deep policy net b) deep value net c) shallow policy net d) shallow value net e) simulation net .



Also, at the leaf, a simulation is performed on a **(circle all that apply)** gpu / cpu . Then, using a fractional weighting of _____ for the call and _____ (these 2 fractions sum to 1.0) for the simulation, the leaf score is backed up the search tree.

5. [4 marks] In MCTS, for a child with w wins and v visits, the function $f(w, v) = (w + t)/(v + 2t)$ is used to measure win rate instead of $g(w, v) = w/v$ because **(circle all that apply)** a) f allows the true value to be estimated more quickly b) f returns a value when v is 0, so never divides by 0 c) f increases the statistical significance of the simulations d) f allows quicker recovery from initial unlucky simulations .

The MCTS UCB1 formula balances the exploitation of a _____ search with the _____ of children which have received fewer _____ than their siblings. For each child j , the formula is **(circle one)** a) $f(w_j, v_j) + c\sqrt{\ln(v_1 + \dots v_t)/v_j}$ b) $f(w, v) + c\sqrt{(v_1 + \dots v_t)/v_j}$ c) $f(w_j, v_j) + c\sqrt{v_j/(v_1 + \dots v_t)}$ d) $f(w, v) + c\sqrt{v_j/\ln(v_1 + \dots v_t)}$.


MCTS can be improved by adding patterns to simulations: eg. in Go, after each simulation move, if a move creates a match with a local (ie. around that move) pattern, then **(circle one)** a) a random move is performed b) the reply move for that pattern is played c) the appropriate player is designated the winner d) the leaf node has its RAVE count increased by 1 .

Eg. in Go, if a white simulation move is as shown in this local 2×2 pattern sequence, then what happens next is _____



6. [3 marks] Before 2000, the strongest Go program was as strong as a human with rank **(circle one)** a) 5 dan b) 9 dan c) 15 kyu d) 30 kyu . MCTS was first used in Go programs around the year _____. Later Clark and Storkey used records from _____ professional games to build a deep _____ neural net that predicts the most popular move correctly with probability about _____. Eg. their net predicts that the first move on the empty _____ x _____ board will be at the **(circle one)** a) 4 4 point b) 5 5 point c) 6 6 point d) 7 7 point. Later the company _____ (owned by Google) wrote AlphaGo, which is about as strong a human with rank **(circle one)** a) 5 dan b) 9 dan c) 15 kyu d) 30 kyu .

- [3 marks] For this Hex position, after black plays at b4, black has a winning virtual connection using cells $\{a5, b5\}$ and $\{a1, a2, a3, b1, b2, c1, c2, d1\}$. Similarly, after black plays at c2, black has a winning virtual connection using cells $\{c1, d1\}$ and $\{a4, a5, b4\}$, and $\{e2, d3\}$ and $\{e4, d5, e5\}$. So, for this position with white to play, white must play at one of the cells in { _____ }, otherwise black can win. For this position with white to play, the set of all winning white moves is { _____ }. So far, the largest Hex boardsize on which winning opening moves have been found is _____ .



10. [2 marks] For the nim state with piles 10 9 8 5, list all winning moves below (if there is no winning move, leave the blank empty). **On the side, show your work.**

From the 10 pile, remove _____ stones.

From the 9 pile, remove _____ stones.

From the 8 pile, remove _____ stones.

From the 5 pile, remove _____ stones.

11. [2 marks] $\begin{matrix} 5 & 2 & 3 \\ 1 & 8 & 4 \\ 7 & 6 & _ \end{matrix}$ The number of inversions of this sliding tile puzzle is _____ (an _____ number) and the number of columns is an odd number, so this puzzle (**circle one**) is / is not solvable.

Consider Python implementations of these algorithms that solve 5x5 (and smaller) sliding tile puzzles: the A* algorithm described in class, and a special-purpose (SP) algorithm using the method from the youtube video discussed in class. (**circle all that apply**)

- a) A* with the Manhattan heuristic is usually faster than A* with the misplaced tiles heuristic.
- b) A* with the Manhattan heuristic is usually faster than SP, because it finds a shortest path through the state space.
- c) A* with the Manhattan heuristic is usually slower than SP, because SP does not search the whole state space.
- d) The runtime for SP on 5x5 puzzles is about 1.25 time the runtime for SP on 4x4 puzzles.

12. [1 marks] Recall: in Go a group of stones is *unconditionally safe* if the opponent cannot kill the group, even if the player always passes. The simplest kind of unconditionally safe group is one that has at least _____ .

13. [6 marks] Consider these games: hex on an 6×6 board, go on a 1×18 board, and tic-tac-toe on an 6×6 board (where to win, you need 4 in a row). For each game you want to write a computer solver (so, an agent that finds a move with best minimax score).

For 6×6 hex, **(circle all that apply)**

a) the game can be solved in a reasonable amount of time using only alphabeta search, since the game tree has only $36! \approx 4 \times 10^{41}$ nodes

b) implementing a transposition table is not difficult, especially since there are no draws

c) the game can be solved in a reasonable amount of time using only alphabeta search, a transposition table, and symmetry pruning, since the solving dag then has about 10^{30} states

d) the game can be solved in a reasonable amount of time using alphabeta search, a transposition table, and pruning with symmetry, mustplay, and inferior cells, since the solving dag then has about 10^{15} states

For 1×18 go, **(circle all that apply)**

a) the game can be solved in a reasonable amount of time using only alphabeta search since the game tree has only $18! \approx 6 \times 10^{15}$ nodes

b) implementing a transposition table is not difficult, since the winner scores at most 18 points

c) the game can be solved in a reasonable amount of time using only alphabeta search and a transposition table since the solving dag then has about 10^{12} states

d) the solver can be improved by recognizing commonly occurring unconditionally safe groups .

For 6×6 tic-tac-toe, **(circle all that apply)**

a) the game can be solved in a reasonable amount of time using only alphabeta search since the game tree has only $36!$ nodes

b) implementing a transposition table is not difficult, as there are only 3 possible outcome values,

c) the game can be solved in a reasonable amount of time using only alphabeta search and a transposition table, since the game is likely to end in a draw,

d) the solver can be improved using mustplay pruning (ie. play here or lose) .