0. On page 0, in the bubbles, write your ***** CCID ***** .
   On pages 0, 1, 2, 3, write your first name, last name and student id.

1. [2 marks] For a 3×3 hex game, assume white plays first.
   Circle all winning first moves.
   a1 a2 a3 b1 b2 b3 c1 c2 c3

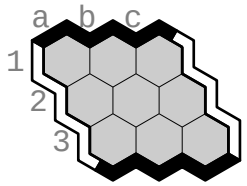

2. [3+3 marks] In github repo program `mcts1.py`, lines 2a,2b,2c below show three different ways to compute the uct value. For each choice of line 2, on my office desktop machine (used in my online lectures), I ran 20 trials of mcts on a 4×4 hex board for 1.0 second each time.

```
1   mean_res = child.results / child.sims
2a  uct = mean_res
2b  uct = mean_res+(self.c*sqrt(log(self.root_node.sims)/child.sims))
2c  uct = mean_res+(self.c*(self.root_node.sims/(child.sims+self.root_node.sims)))
```

   a) For each version, circle the number closest to the average number of simulations in each trial.

   version 2a: 110000 130000 160000 1100 1300 1600 11000 13000 16000

   version 2b: 110000 130000 160000 1100 1300 1600 11000 13000 16000

   version 2c: 110000 130000 160000 1100 1300 1600 11000 13000 16000

   Justify your answer briefly:

   b) For each version, out of 20 trials, give the total number of times the move returned by mcts was a winning move.

   version 2a _____ /20 version 2b _____ /20 version 2c _____ /20

   Justify your answer briefly:

3. [1+1+4+2 marks] This is a winning 3×3 hex strategy:

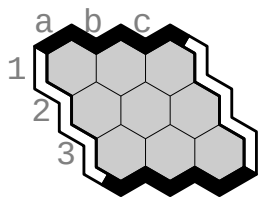

$a3 \wedge (b3 \wedge (c2 \vee c3) \ \vee \ c1 \wedge (b2 \vee b1 \wedge (a1 \vee a2)))$

a) is this a strategy for black or for white or for both?

b) is this a first-player or second-player strategy?

c) explain the strategy in words

d) If you represented the strategy using a dag, how many nodes would be in the dag? Justify.

4. Here is output from an execution of the class github mcts on this hex board. Recall cells 5,6,9,10 are cells a1,b1,a2,b2. The diagram shows the search tree after sim 1. At each node, label x, y is wins,sims. Continue the diagram: show what it looks like after sim 6 (so, after all sims have executed).

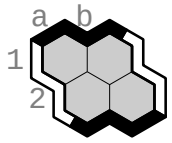


```
root expand * > 5, 6, 9, 10, done
sim  1.    *  6 roll  9  5 10 parent loss
sim  2.    *  9 roll 10  6 parent win
sim  3.    *  5 roll 10  9 parent win
sim  4.    * 10 roll  6  5  9 parent loss
expand * 5 > 6, 9, 10
sim  5.    *  5  9 roll 10  6 parent win
expand * 9 > 5, 6, 10
sim  6.    *  9  5 roll 10  6 parent win
```

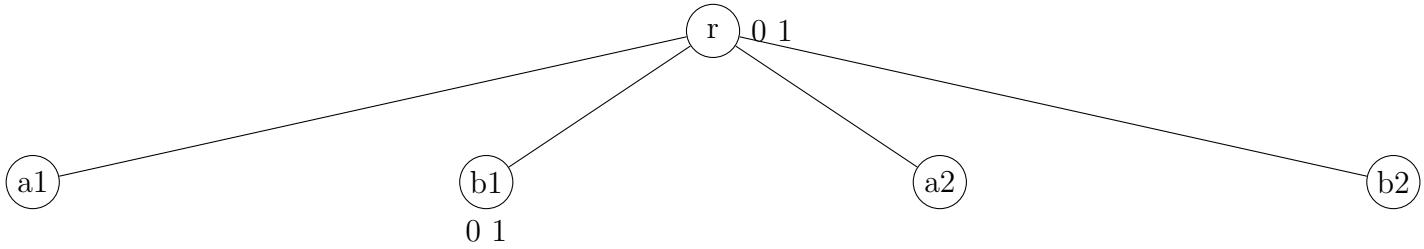

0. On page 0, in the bubbles, write your ***** CCID ***** .
   On pages 0, 1, 2, 3, write your first name, last name and student id.

1. [2 marks] For a 3×3 hex game, assume white plays first.
   Circle all winning first moves.
   a1 b1 c1 a2 b2 c2 a3 b3 c3

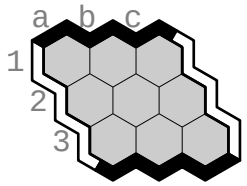

2. [3+3 marks] In github repo program `mcts1.py`, lines 2a,2b,2c below show three different ways to compute the uct value. For each choice of line 2, on my office desktop machine (used in my online lectures), I ran 20 trials of mcts on a 4×4 hex board for 1.0 second each time.

```
1   mean_res = child.results / child.sims
2a  uct = mean_res+(self.c*sqrt(log(self.root_node.sims)/child.sims))
2b  uct = mean_res+(self.c*(self.root_node.sims/(child.sims+self.root_node.sims)))
2c  uct = mean_res
```

   a) For each version, circle the number closest to the average number of simulations in each trial.

   version 2a: 16000 11000 13000 1600 1100 1300 160000 110000 130000

   version 2b: 16000 11000 13000 1600 1100 1300 160000 110000 130000

   version 2c: 16000 11000 13000 1600 1100 1300 160000 110000 130000

   Justify your answer briefly:

   b) For each version, out of 20 trials, give the total number of times the move returned by mcts was a winning move.

   version 2a _____ /20 version 2b _____ /20 version 2c _____ /20

   Justify your answer briefly:

3. [1+1+4+2 marks] This is a winning 3×3 hex strategy:

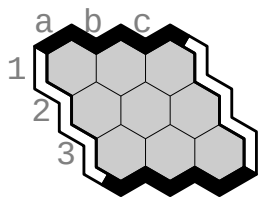

$c1 \wedge (c2 \wedge (b3 \vee c3) \ \vee \ a3 \wedge (b2 \vee a2 \wedge (a1 \vee b1)))$

a) is this a strategy for black or for white or for both?

b) is this a first-player or second-player strategy?

c) explain the strategy in words

d) If you represented the strategy using a dag, how many nodes would be in the dag? Justify.

4. Here is output from an execution of the class github mcts on this hex board. Recall cells 5,6,9,10 are cells a1,b1,a2,b2. The diagram shows the search tree after sim 1. At each node, label x, y is wins,sims. Continue the diagram: show what it looks like after sim 6 (so, after all sims have executed).

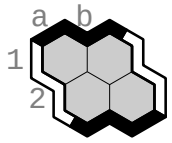


```
root expand * > 5, 6, 9, 10, done
sim  1.    *  9 roll  5 10  6 parent loss
sim  2.    *  6 roll  5 10 parent win
sim  3.    *  5 roll  9 10  6 parent loss
sim  4.    * 10 roll  5  6 parent win
expand * 6 > 5, 9, 10
sim  5.    *  6 10 roll  9 parent loss
expand * 10 > 5, 6, 9
sim  6.    * 10  9 roll  6 parent loss
```

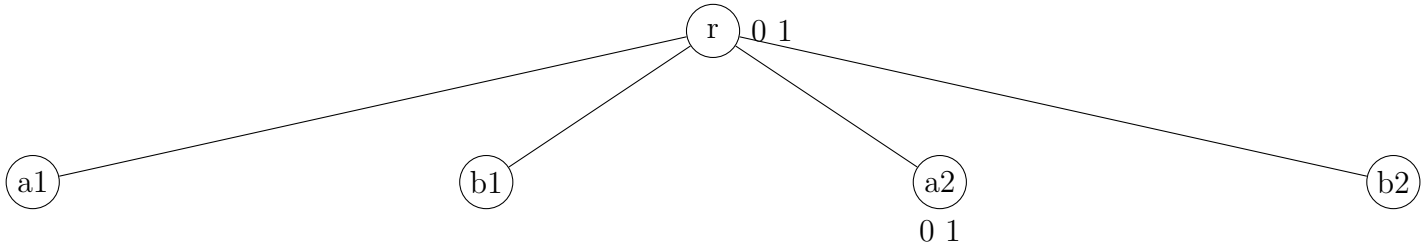

0. On page 0, in the bubbles, write your ***** CCID ***** .
   On pages 0, 1, 2, 3, write your first name, last name and student id.

1. [2 marks] For a 3×3 hex game, assume black plays first.
   Circle all winning first moves.
   a1 a2 a3 b1 b2 b3 c1 c2 c3

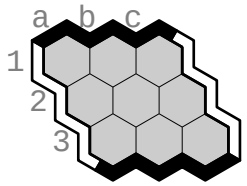

2. [3+3 marks] In github repo program `mcts1.py`, lines 2a,2b,2c below show three different ways to compute the uct value. For each choice of line 2, on my office desktop machine (used in my online lectures), I ran 20 trials of mcts on a 4×4 hex board for 1.0 second each time.

```
1   mean_res = child.results / child.sims
2a  uct = mean_res+(self.c*(self.root_node.sims/(child.sims+self.root_node.sims)))
2b  uct = mean_res
2c  uct = mean_res+(self.c*sqrt(log(self.root_node.sims)/child.sims))
```

   a) For each version, circle the number closest to the average number of simulations in each trial.

version 2a:	1300	1600	1100	13000	16000	11000	130000	160000	110000
version 2b:	1300	1600	1100	13000	16000	11000	130000	160000	110000
version 2c:	1300	1600	1100	13000	16000	11000	130000	160000	110000

   Justify your answer briefly:

   b) For each version, out of 20 trials, give the total number of times the move returned by mcts was a winning move.

   version 2a _____ /20 version 2b _____ /20 version 2c _____ /20

   Justify your answer briefly:

3. [1+1+4+2 marks] This is a winning 3×3 hex strategy:

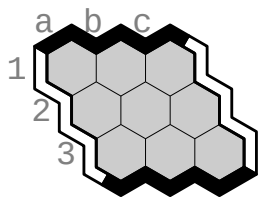

$c1 \wedge (b1 \wedge (a1 \vee a2) \ \vee \ a3 \wedge (b2 \vee b3 \wedge (c2 \vee c3)))$

a) is this a strategy for black or for white or for both?

b) is this a first-player or second-player strategy?

c) explain the strategy in words

d) If you represented the strategy using a dag, how many nodes would be in the dag? Justify.

4. Here is output from an execution of the class github mcts on this hex board. Recall cells 5,6,9,10 are cells a1,b1,a2,b2. The diagram shows the search tree after sim 1. At each node, label x, y is wins,sims. Continue the diagram: show what it looks like after sim 6 (so, after all sims have executed).

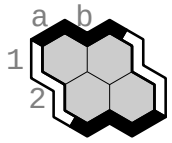


```
root expand * > 5, 6, 9, 10, done
sim  1.    * 10 roll  5  6 parent win
sim  2.    *  5 roll  9 10  6 parent loss
sim  3.    *  9 roll  5  6 parent win
sim  4.    *  6 roll  9 10 parent win
expand *  6 > 5, 9, 10
sim  5.    *  6 10 roll  5  9 parent win
expand *  9 > 5, 6, 10
sim  6.    *  9  6 roll  5 parent loss
```