

5COSC023W - Tutorial 4/5/6 Exercises

1 A Tic-Tac-Toe Game

In this exercise we will develop a Android version of the classic tic-tac-toe game which self-learns how to play as expert by playing games against itself.

This is a machine learning approach, specifically from the area of reinforcement learning where an agent learns a task by interacting with the environment. **(if you would like to tackle machine learning as part of your final year project and under my supervision, we could have a chat at some point).**

The game is played with two players, **X** (user) and **O** (computer), who take turns marking the spaces in a 3×3 grid. The **X** player goes first. The player who succeeds in placing three respective marks in a horizontal, vertical, or diagonal row wins the game (see Figure 1 below).

X	X	O
	X	O
	O	X

Figure 1: An instance of the tic-tac-toe board where the X player has won since he managed to place three 'X' in a diagonal.

1. Implement an Android version of the game that a human player can use to play against the computer. Initially, the computer player chooses a random cell as its next move.
2. Implement an intelligent version of the computer player (i.e. replacing the `RandomPlayer` class) by extending class `Player`. The intelligent computer player is able to defend itself, i.e. if the human player has already placed 2 X in a row, column or diagonal, the computer player will choose the free slot which will prevent the human to win in their next move.
3. Extend your intelligent player class so that it is able to choose winning positions, i.e. if the computer player has already placed 2 O in a row, column or diagonal then it places the next O in the slot which completes 3 O to win the game.

If there is no winning move, but there is a defending move, the intelligent computer player will choose the defending one. If there is neither a winning or defending move, the intelligent player will choose a random valid move (i.e. a random slot among all empty slots).

4. **This part will be covered in Week 5** - Implement a *Computer Player* which self learns to play, by playing hundreds of games against itself. The “trained” computer player is then placed to put against a human.

Can you beat the computer?