

AI Tutorials

Machine Learning


Q1. [Decision Tree] We want to predict the outcome of the next cricket match between the two top-ranked teams: India and Australia. We are given the following dataset from recent matches. The outcome of the matches are T if India wins and F if Australia wins.

Time	Match Type	Pitch Type	Outcome
Evening	T20	Fast	T
Day and Night	ODI	Dusty	T
Day	Test Match	Bouncy	T
Day and Night	Test Match	Neutral	F
Day and Night	T20	Dusty	F
Day and Night	ODI	Fast	T
Day and Night	ODI	Bouncy	T
Day and Night	ODI	Bouncy	T
Evening	T20	Fast	T
Day and Night	ODI	Dusty	F
Day	Test Match	Bouncy	F
Day	T20	Neutral	F
Day and Night	T20	Dusty	F
Day and Night	T20	Fast	T
Day and Night	ODI	Bouncy	T
Day and Night	ODI	Dusty	T

You are required to create a Decision Tree from this data and use it to decide whether India is likely to win the next match against Australia.

- What is the entropy of Outcome? Write the formula for entropy and then use it in your computation.
- Which attribute should you choose as the root of a decision tree?
- What is the information gain of the attribute you chose in the previous question? Write the formula and then use it in your computation.
- Build the complete decision tree.
- We are given that the next match is an ODI to be played as a Day and Night match and the pitch is dusty. Use the decision tree to decide whether India is likely to win the match. If the decision is not certain, then indicate the probability of win.

Q2. [Neural Network] Take the following Neural Network.

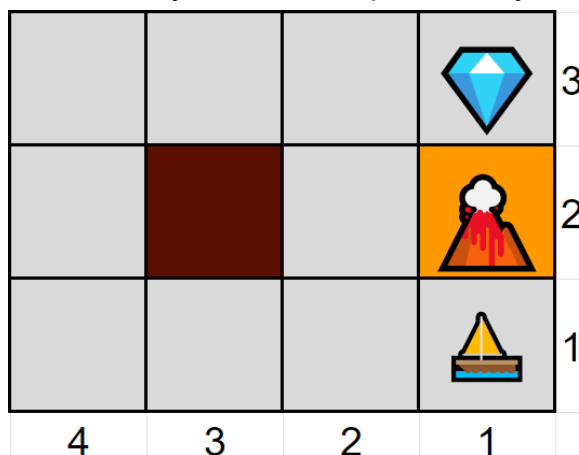


Ix : Inputs; Bx : Biases; Hx : Hidden Layers; Ox : Outputs.

If all the Activation Functions are ReLU, then answer the following questions.

- What will be the output for Input [1, 1]?
- What will be the total loss for the same input if the actual output is [1, 0] and the loss function is Mean Square Error (MSE)?
- What will be the output and loss if sigmoid activation functions are used in the output layer instead of ReLU?

Q3. [Reinforcement Learning] Imagine you are a treasure hunter, but you do not collect the treasure; you draw treasure maps for your descendants. The Grid World shown below is the overall map of the island you are currently exploring. You don't have any prior information about where the treasure is, and due to smoke, you cannot see what is in the next block. You cannot go to the stone block, and if you fall into the volcano, you will be dead, but you have a magic power to respawn in your boat. You can take four actions **[Up, Right, Down, Left]** to move from one block to the next if a block is available in that direction. Otherwise, you will end up in that block only. The treasure is worth 100 diamonds (points), but if you're dead, then you also lose 100 (points). Each time you take a step, it costs you one point (living cost).




If you use Q-Learning with ϵ -Greedy Policy to draw the Map, please answer the following questions:

[NB: Initial Q values are zero. Please follow the action order [Up, Right, Down, Left] to break any tie. $\epsilon = 0.1$, α (Learning Rate) = 0.1 and γ (discount factor) = 0.9]

- a. Please show the Map (q-Values) after 2-Iteration, if you do not take any ϵ actions. (10 Step Max)

For example, Map after 1 Iteration.



- b. From that can you guess what is going to be the Optimal path to treasure?