

Consider the game of tic-tac-toe. The game board is a 3x3 grid, all nine cells of which are empty in the beginning. Two players, 1 and 2, take turns to mark the cells of the board one by one. The game is won by the player who manages to mark all three cells in a row or column or across any of the diagonals.

Consider the flattening of the game board grid into a vector, the first three elements of which denotes the top-row, the next three elements denote the middle row, and the last three elements denote the last row. Like this: [(1, 1), (1, 2), (1, 3), (2, 1), (2, 2), (2, 3), (3, 1), (3, 2), (3, 3)].

A configuration in the game board is a vector of nine values (taken from 0, 1, and 2 – to denote that the corresponding cell is empty, marked by player 1, and marked by player 2, respectively). For example, a configuration may look like this: [1, 0, 2, 0, 1, 2, 0, 0, 0].

Let us consider only those configurations where the number of 1's and 2's are equal, both either two or three, and there's at least one unmarked cell (i.e., a cell with a 0). Let us call these as *valid* configurations. Let us also assume that a valid configuration is not already winning for either of the players.

We want to train a network to predict a *good* move for player 1, given a valid configuration. A *move* for player i here would simply mean that one of the 0's in the vector is replaced by i . We call such a move *bad* if:

- there's a way for player 1 to win at this stage and it does not make a winning move, *or*
- there's no way for player 1 to win at this stage; however, player 1's move allows player 2 to win from the resulting configuration in one move.

A move which is not *bad* is classified as *good*. Your task is to do the following.

1. Create data (training and testing) for this using z3. Essentially, you need to generate a tuple (config, index-for-player-1-move), and its classification as good or bad. For example:

([1, 0, 2, 0, 1, 2, 0, 0, 0], 8) → good
 ([1, 0, 2, 0, 1, 2, 0, 0, 0], 7) → bad
 ([1, 0, 2, 0, 1, 2, 0, 0, 0], 1) → bad
 ([1, 1, 2, 0, 0, 2, 0, 0, 0], 8) → good
 ([1, 1, 2, 0, 0, 2, 0, 0, 0], 4) → bad
 ([1, 1, 2, 0, 0, 2, 0, 0, 0], 6) → bad

Notice, however, that the marking of 1 and 2 is symbolic. So, it might be better to use the following encoding for symbols:

0 → 100, 1 → 010, 2 → 001

and, similarly, 9 bits to encode the index corresponding to the cell being marked by player 1. Like this:

8 → 000000001, 7 → 000000010, 6 → 000000100
 5 → 000001000, 4 → 000010000, 3 → 000100000
 2 → 001000000, 1 → 010000000, 0 → 100000000

So, instead of using the input vector of 10 integers (first 9 for the board, and the last one for the index), we will have an input vector of 36 bits (3 bits each for 9 numbers for the board, and 9 bits for the index). The output (1 for *good* move, 2 for *bad* move) can similarly be encoded using two bits.

The classification – *good* or *bad*, for a valid configuration – must be found by checking (using z3) the two conditions, stated earlier, in which a move is classified as bad.

2. Create a network and train it using this data. Select (randomly) 50–60% of the generated data for training, and use the remaining data for testing. You may manipulate the dimensions of this network to get an accuracy of at least 85% on the training data, and at least 75% on the testing data. Please keep at least 2 hidden layers, and a total of 20 hidden neurons or more. Use ReLU as the activation function, except in the last layer where you may not use any activation function, or use softmax. However, while encoding this for property checking, you should not encode softmax. Instead, you should simply check which output value is bigger.
3. In the trained network, you should check the following correctness property – *Does there exist an input (a valid configuration, and an index for player 1's move) which is classified as good by the network, but from the configuration resulting from it, it is possible for player 2 to make one move and win?* Use both z3 and Marabou to produce such an input, if one exists, or to show lack of one. Compare the time taken by the two tools.
4. In the trained network, also check the following property – *Does there exist an input, where the configuration has the center cell and at least one of the corner cells marked by player 1, which is classified as good by the network, but from the configuration resulting from it, it is possible for player 2 to make one move and win?* Once again, you should use both z3 and Marabou to produce such an input, if one exists, or to show lack of one. And compare the time taken by the two tools.

Note that you can verify other, more involved, properties as well. For instance, you may want to know if it is possible to start from a valid board where both the players have marked two cells each, pick a *good* move for player 1, and then assuming that player 2 plays anywhere (but does not win), you once again pick a *good* move for player 1, but it is still possible for player 2 to make one move and win. This is not a part of the assignment, but I encourage you to try such properties.