

LIFE AND WELL-BEING SCIENCE

Create your own AI

ALEXANDER FARRUGIA

People may think that making your own artificial intelligence (AI) is difficult or technical. However, for certain tasks like playing a simple game, it is not hard to build your own self-learning AI. What's more, you do not even need a computer or electronic device to do this!

A simple yet challenging game is the two-player game called Nim. In Nim, there are counters arranged in at least two piles. At each turn, a player chooses one pile and takes as many counters as they want from it. The player taking the last counter loses the game. The objective of the game, then, is to force your opponent to take the last counter.

For example, suppose our Nim game has two piles, Pile A, containing four counters, and Pile B, containing three counters. A game may go like this: Player 1, starting first, takes one counter from Pile A. Player 2 takes two counters from Pile A. Player 1 takes all three counters from Pile B.



The 13 possible states of our example Nim game, represented as 13 cups, each of which contain folded papers representing all the possible playable moves from that game state.

Now Player 2 must take the last remaining counter from Pile A, losing the game.

To create an AI that plays Nim, one first needs to figure out all its distinct possible states. In the game described above, there are 13 different states; thus, we require 13 containers, for example paper cups, each representing a distinct game state. Each cup would then contain every possible legal move that can be made from that particular game state; these can be represented by folded pieces of

paper, each depicting a legal move that can be made at that game state.

When it is the AI's turn to play, a paper is drawn at random from the cup representing the game's current state, and the move depicted on that piece of paper is played. If the AI eventually wins, all the papers drawn are folded and put back in their respective cups. If the AI loses, however, all the papers except the one representing the last move are put back in their cup. However, if by doing so, a cup remains

empty, then that paper is still put back in its cup, and the paper representing the game's previous move is discarded instead.

It is a fascinating experience witnessing a 'device' that has no knowledge of the strategies of a game, eventually learning from its mistakes until it masters that game. I invite you to create your own AI and to experience this yourselves!

Alexander Farrugia is a mathematics lecturer at the University of Malta Junior College.

MYTH DEBUNKED

Is AI a recent academic discipline?

ALEXANDER FARRUGIA

Contrary to what people may think, AI has been around for decades.

The April 1984 edition of *ANTIC* magazine (a copy of which may be found at <https://www.atarimagazines.com/index/index.php?issue=v3n1>) featured an ATARI computer programme, written by Joseph Hafner, that implemented a self-learning AI playing Tic-Tac-Toe.


This AI, originally devised by Donald Michie in 1961, was called *MENACE*, short for 'Matchbox Educable Noughts and Crosses Engine'. *MENACE* used 304 matchboxes to represent every state of the game. Each matchbox contained differently coloured beads that were drawn at random to represent the move that the AI should make. The AI learned the best moves by being encouraged to replay winning moves and avoid losing moves – this technique is called reinforcement learning.

This shows that AI has been around for quite a long time; indeed, AI was founded as an academic discipline in 1956. Moreover, in 1950, Alan Turing devised a test for machine intelligence, aptly called the Turing test – that's even before AI was officially a field of study!

Incidentally, in 2015, Matt Parker and his team recreated the original *MENACE*, with matchboxes, beads and all, at the Manchester Science Festival. After losing a considerable number of games, the AI learned the correct strategy and soon after became unbeatable.

A video of this event may be found at https://youtu.be/R9c_neaxeU, where many details are elucidated.

PHOTO OF THE WEEK



A screenshot of the ATARI computer programme 'Matchbox Tic-Tac-Toe', written in 1984 for ATARI computers by Joseph Hafner and published by *ANTIC* magazine. Despite a dismal record of two wins, four draws and 25 losses, the artificial intelligence was starting to learn how to play correctly, as evidenced by the draw achieved in the 32nd game. For the author of this page, this game represents his first encounter with a self-learning AI – he used to play this game when he was a child!

DID YOU KNOW?

- The game of Nim is solved, that is, the best strategy for each player is known for any number of piles and for any number of counters in each pile.
- There are several variants of Nim. In one variant, the player taking the last counter wins the game, instead of loses the game. In another variant, each player must take counters from the pile containing the largest number of counters in that turn.
- Chess is not a solved game. Although most chess experts agree that a chess game would end in a draw if both players play the best possible move at every turn, this has not been proven mathematically yet. However, chess is solved for positions having up to seven pieces.
- Game-playing AIs aid in discovering innovations in other fields of AI, especially those utilising reinforcement learning, such as natural language processing and robotics.

For more trivia, see www.um.edu.mt/think

SOUND BITES

- In 2017, DeepMind, an AI research lab serving as a subsidiary of Google, released a paper preprint describing AlphaZero, an AI that learns the strategies of three games (chess, Shogi and Go) by playing against

itself. After training for a maximum of 34 hours, AlphaZero managed to comprehensively beat the best AI for each game at that point in time – Stockfish, Elmo and AlphaGo respectively. A newer version of AlphaZero was released by DeepMind in 2022 that discovered a faster way to perform matrix multiplication. Also in 2022,

another programme, muZero, was released by DeepMind that plays against itself to learn how to play chess, Shogi, Go and ATARI 2600 games without even knowing the rules of any of these games!

For more science news, listen to Radio Mocha on www.fb.com/RadioMochaMalta/

CONTRIBUTORS

EDWARD DUCA
JEAN-PAUL EBEJER
DANIELLE MARTINE FARRUGIA
JOSEF BORG
CLAUDIA BORG
ALEXANDER FARRUGIA
MOHAMMED DAOUD

E-MAIL

SCI-SUNDAY@UM.EDU.MT