Introduction to Reinforcement Learning

Reinforcement Learning (RL) stands as a significant branch of artificial intelligence, distinguished by its emphasis on enabling agents to make decisions through iterative interactions with their environment. Unlike traditional supervised learning approaches, where explicit input-output pairs guide model training, RL engages agents in autonomous exploration, learning from the consequences of their actions.

In the context of artificial intelligence, an "agent" is a computer or system that attempts to play the role of a "user" in the program. The term "agent" encompasses a broad range of entities, from simple programs executing predefined actions to more complex systems employing learning algorithms to adapt and improve their decision-making over time. The key components of an agent typically include

- **perception**, which involves sensing and interpreting information from the environment
- **decision-making**, where the agent selects actions based on its goals and current state
- **action**, the execution of chosen actions to impact the environment.

Agents play a crucial role in various AI applications, from game-playing algorithms and autonomous vehicles to virtual assistants and industrial automation systems. Understanding and modelling agents is essential for developing intelligent systems that can effectively navigate and respond to dynamic and uncertain environments.

One version of RL is the "Quality-learning" or "Q-Learning" algorithm, a fundamental technique for training agents in decision-making tasks.

Q-learning operates by **estimating the value of taking specific actions** in distinct states and **updating a table based on the rewards observed** during the learning

process. Repeating this process (choose an action, observe rewards) many times allows agents to develop optimal strategies over time, aligning their actions with the goal of maximizing cumulative rewards.

The RL framework comprises several essential components, including the agent, environment, state, action, and reward. The agent, tasked with decision-making, interacts with an environment defined by various states. The agent takes actions in these states, and the environment responds with rewards or penalties, providing the feedback (or "reinforcement") necessary for the agent to adjust its strategy.

Exploring reinforcement learning concepts is particularly insightful when applied to gaming scenarios. Consider the classic game of Tic-Tac-Toe, a simple yet illustrative environment for understanding reinforcement learning principles.

In the context of Tic-Tac-Toe, we have:

Agent - A computer program learning to play the game of Tic-Tac-Toe.

Environment - A game of Tic-Tac-Toe

States - various configurations of the Tic-Tac-Toe board as the game is played

Actions - the decision by the agent about where to place their token in a given board configuration.

Rewards - Moves in the game that directly result in the agent winning the game, or tying the game are rewarded. The reward for winning is greater than the reward for Tying the game, because this is the preferred outcome.