What is Deep Reinforcement Learning?

Types of Machine Learning





Reinforcement Learning

- **Reinforcement learning** (**RL**) is an area of <u>machine</u> <u>learning</u> concerned with how <u>software agents</u> ought to take <u>actions</u> in an environment so as to maximize some notion of cumulative reward. Reinforcement learning is one of three basic machine learning paradigms, alongside <u>supervised</u> <u>learning</u> and <u>unsupervised learning</u>.
- It differs from supervised learning in that labelled input/output pairs need not be presented, and sub-optimal actions need not be explicitly corrected. Instead the focus is finding a balance between exploration (of uncharted territory) and exploitation (of current knowledge)

Reinforcement Learning



Some key terms that describe the basic elements of an RL problem are:

- Environment Physical world in which the agent operates
- **State** Current situation of the agent
- **Reward** Feedback from the environment
- Policy Method to map agent's state to actions
- Value Future reward that an agent would receive by taking an action in a particular state

- An RL problem can be best explained through games.
- Let's take the game of <u>PacMan</u> where the goal of the **agent** (PacMan) is to eat the food in the grid while avoiding the ghosts on its way.
- The grid world is the interactive **environment** for the agent where it acts.
- Agent receives a **reward** for eating food and punishment if it gets killed by the ghost (loses the game).
- The **states** are the location of the agent in the grid world and the total cumulative reward is the agent winning the game.
- In order to build an optimal **policy**, the agent faces the dilemma of exploring new states while maximizing its overall reward at the same time. the best overall strategy may involve short term sacrifices. Therefore, the agent should collect enough information to make the best overall decision in the future.

What are some of the most used Reinforcement Learning algorithms?

- **Q-learning** and **SARSA** (State-Reward-State-Action) are two commonly used model-free RL algorithms.
- Deep Deterministic Policy Gradient(DDPG) is a model-fr policy, actor-critic algorithm th tackles this problem by learnir policies in high dimensional, continuous action spaces.

What are the practical applications of Reinforcement Learning?

- <u>AlphaGo</u>
- Robotic Manipulation