

Reinforcement Learning (RL)

By Aylin Muzaffarli

What is Reinforcement Learning?

Reinforcement Learning (RL) is a branch of machine learning where an **agent** learns to make decisions by interacting with an **environment**. The agent performs **actions**, observes the **state** of the environment, and receives **rewards** as feedback. The goal is to learn a **policy** that maximizes cumulative rewards over time.

In simple terms: RL is learning by trial and error, guided by rewards.

Core Concepts

- **Agent:** The decision-maker (e.g., a robot, a program, a game character).
 - **Environment:** The world the agent interacts with (e.g., a game grid, a simulation).
 - **State:** The current situation of the environment.
 - **Action:** Any move or decision the agent can make.
 - **Reward:** Feedback received after an action. Positive rewards encourage behavior; negative rewards discourage it.
 - **Policy:** The strategy the agent uses to decide which actions to take in each state.
-

Why RL is Important

RL is the backbone of many cutting-edge AI systems:

- Game-playing AI (like AlphaGo)

- Robotics (learning to walk, pick, or manipulate objects)
 - Recommendation systems (deciding what content to show)
 - Autonomous vehicles (deciding how to drive safely)
-

Beginner-Level Insights

- RL is different from supervised learning: you don't give it the "right answer" for every step; the agent **learns from feedback**.
 - Rewards guide the learning process; designing a good reward system is crucial.
 - RL can be applied to small simulations to **understand patterns** before scaling up to complex systems.
-

Recommended Resource

- [Introduction to Reinforcement Learning – University of Verona PDF](#) – concise, beginner-friendly theory.