

Using this [algorithm](#), the [machine](#) is trained to make specific [decisions](#). It works this way: the machine is exposed to an [environment](#) where it trains itself continually using [trial](#) and [error](#). This machine learns from past [experience](#) and tries to capture the best possible [knowledge](#) to make accurate business decisions. Example of Reinforcement Learning: Markov Decision Process.

Machine learning [algorithms](#) can be divided into 3 broad categories—[supervised learning](#), [unsupervised learning](#), and [reinforcement learning](#). Supervised learning is useful in cases where a property (label) is available for a certain [dataset](#) (training set), but is missing and needs to be predicted for other instances. Unsupervised learning is useful in cases where the challenge is to discover implicit relationships in a given unlabeled dataset (items are not pre-assigned). Reinforcement learning falls between these 2 extremes—there is some form of feedback available for each predictive step or action, but no precise label or error message.

From:

<https://neurosurgerywiki.com/wiki/> - **Neurosurgery Wiki**

Permanent link:

https://neurosurgerywiki.com/wiki/doku.php?id=reinforcement_learning

Last update: **2024/06/07 02:58**

