Machine Learning: Reinforcement Learning

By Gunnar Smith

What is Machine Learning?

Machine learning is the use of data and algorithms to imitate intelligent human behavior

3 types of Machine Learning

1. Supervised learning
    a. The machine is trained on a statistically representative set of example inputs and corresponding outputs.

1. Unsupervised learning
    a. The machine is not provided labeled examples or patterns in which the machine must uncover patterns without the correct answer

1. Reinforcement learning
    a. The Machine learns by continuously improving its model based on feedback from experiences.

Reinforcement Learning Terms

Value function: the algorithm to determine the value of being in a state, the probability of receiving a future reward

Agent: The entity which performs action in an environment.

Policy: The strategy applied by the agent

State: Refers to the current situation returned by the environment after agent's action

Value: The expected long-term return compared to the short-term reward

Environment: The scenario that an agent has to face

Reward: A return value given to the agent after it performs a task

Source | www.guru99.com/reinforcement-learning-tutorial

Reinforcement Learning Types

Model-Based

Model-based learning algorithms attempt to model the environment where the agent choose the optimal policy based on it's learned model

I.E Learn the effects of actions

Model-Free

Model-free learning algorithms rely on the agent's experience through trial-and-error for setting up the optimal policy

I.E Learn the values of actions

Source | www.guru99.com/reinforcement-learning-tutorial

Model-Based Reinforcement learning

Objective: an agent trying to understand its environment and creating a model for it based on its interactions with this environment

Model preference

In this system, preferences take priority over the consequences of the action

The agent will always try perform an action that will maximize reward regardless what the action may cause

Model-Free Reinforcement Learning

Objective: seek to learn the consequences of their actions through experience via algorithms such as Policy Iteration and Value iteration

Policy iteration

Evaluates a policy and then uses these values to improve that policy. Policy iteration tries to maximize the value of the next state

EX. Video game rewards

Value iteration

The agent will find a optimal policy that maximizes the value function

EX. Automatic parking policies

Machine Learning: Reinforcement Learning Applications

- Autonomous Cars
  - Trajectory optimization
  - Motion Planning
  - Dynamic pathing
  - Controller optimization
- Healthcare
  - Dynamic Treatment Regimes (DTRs)
- Personalised Recommendations
  - Ad recommendation system
  - Product recommendation system
- Traffic Light Control
- Video games

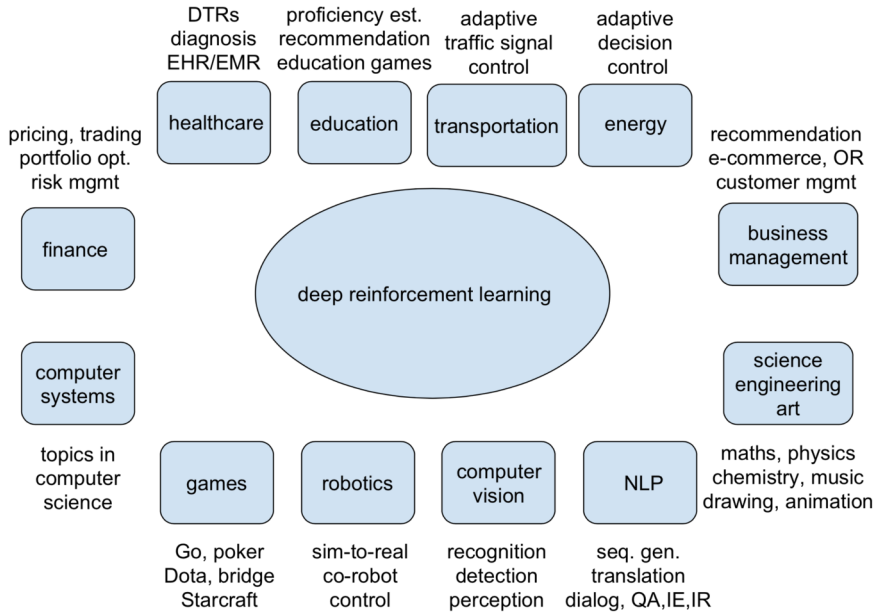

Yuxi Li, Deep Reinforcement Learning, arXiv, 2018

Source | https://www.javatpoint.com/applications-of-machine-learningç

Reinforcement Learning Shortcomings

1. Too much reinforcement learning can lead to and overload of states which diminish the results
2. Reinforcement learning is not preferable when it comes to solving simple problems
3. Parameters may affect legitimacy of results
4. The necessity of a reward function
5. Lack of interpretability

Source | https://www.asquero.com/article/advantages-and-disadvantages-of-different-types-of-machine-learning-algorithms/

Any Questions?

Sources

https://core-robotics.gatech.edu/2021/01/19/bootcamp-summer-2020-week-3-value-iteration-and-q-learning/

https://neptune.ai/blog/model-based-and-model-free-reinforcement-learning-pytennis-case-study

https://www.guavus.com/ai-vs-machine-learing-vs-data-mining-whats-big-difference-part-2/

https://www.ibm.com/cloud/learn/machine-learning

https://pythonistaplanet.com/pros-and-cons-of-reinforcement-learning/

https://www.baeldung.com/cs/ml-value-iteration-vs-policy-iteration