



MACHINE LEARNING: REINFORCEMENT LEARNING^[A]

Dr. A. R. Muralidharan*

M.Sc., M.S.I.T., M.Phil., Ph.D Assistant Professor in Statistics Department of Computational Sciences College of Natural and Computational Sciences Ethiopia.

Article Received on 18/11/2021

Article Revised on 08/12/2021

Article Accepted on 28/12/2021

***Corresponding Author**

Dr. A. R. Muralidharan

M.Sc., M.S.I.T., M.Phil.,
Ph.D Assistant Professor in
Statistics Department of
Computational Sciences
College of Natural and
Computational Sciences
Ethiopia.

ABSTRACT

In this current scenario, the application and domination of machines (computers and related gadgets) are unavoidable around the globe. The machine can be replaced in human brains, such as decision making and prediction of the cause and effects. The word “Machine Learning” is the current trend. All Industries and social domains are turned to this buzzword. Thus it has important in aware of these related techniques. There are several methods to handle the data based on the situations faced by the end-user. This classification includes. Supervised, semi-supervised, unsupervised and reinforcement learning (RL). In this article, the author made the discussion on Reinforcement learning and its related concepts.

KEYWORDS: Machine learning, reinforcement learning, artificial intelligence, deep learning,

INTRODUCTION

Machine learning must be one in all the fastest-growing fields in engineering. it's not only that the information is continuously getting “bigger,” but also the speculation to process it and switch it into knowledge. In various fields of science, from astronomy to biology, but also in daily life, as digital technology increasingly infiltrates our daily existence, as our digital footprint deepens, more data is continuously generated and picked up. Whether scientific or personal, data that just lies dormant passively isn't of any use, and smart people are finding ever new ways to form use of that data and switch it into a useful product or service. during this transformation, machine learning plays a bigger and bigger role.

Machine learning is related to artificial intelligence and deep learning and can be segregated as follows:

- **Artificial Intelligence (AI)** is the broadest term applied to any technique that enables computers to mimic human intelligence using logic, if-then rules, decision trees, and machine learning (including deep learning).
- **Machine Learning** is the subset of AI that includes abstruse statistical techniques that enable machines to improve at tasks with the experience gained while executing the tasks. If we have input data x and want to find the response y , it can be represented by the function $y = f(x)$. Since it is impossible to find the function f , given the data and the response (due to a variety of reasons discussed in this book), we try to approximate f with a function g . The process of trying to arrive at the best approximation to f is through a process known as machine learning.
- **Deep Learning** is a scalable version of machine learning. It tries to expand the possible range of estimated functions. If machine learning can learn, say 1000 models, deep learning allows us to learn, say 10000 models. Although both have infinite spaces, deep learning has a larger viable space due to the math, by exposing multilayered neural networks to vast amounts of data.



Figure 1: Learning data: Traditional and Machine (programming)

Machine Learning (ML) is utilized anywhere from automating mundane tasks to offering intelligent insights, industries in every sector try to take pride in it. you'll already be employing a tool that utilizes it. for example, a wearable fitness tracker like Fitbit, or an intelligent home assistant like Google Home. But there are rather more samples of ML in use.

- **Prediction** — Machine learning can also be used within the prediction systems. In the loan example, to compute the probability of a fault, the system will must classify the available data in groups.

- Image recognition — Machine learning is also used for face detection in a very picture still. There is a separate category for each person in an exceedingly database of several people.
- Speech Recognition — it is the interpretation of spoken words into the text. it's employed in voice searches and more. Voice user interfaces include voice dialing, call routing, and appliance control. it should be used a straightforward data entry and so the preparation of structured documents.
- Medical diagnoses — ML is trained to acknowledge cancerous tissues.
- Financial industry and trading — companies use ML in fraud investigations and credit checks.

Classification of Machine Learning

Machine learning implementations are classified into three major categories, depending on the nature of the learning “signal” or “response” available to a learning system which is as follows:-

- 1. Supervised learning:** When an algorithm learns from example data and associated target responses that can consist of numeric values or string labels, such as classes or tags, in order to later predict the correct response when posed with new examples comes under the category of Supervised learning. This approach is indeed similar to human learning under the supervision of a teacher. The teacher provides good examples for the student to memorize, and the student then derives general rules from these specific examples.
- 2. Unsupervised learning:** Whereas when an algorithm learns from plain examples without any associated response, leaving to the algorithm to determine the data patterns on its own. This type of algorithm tends to restructure the data into something else, such as new features that may represent a class or a new series of un-correlated values. They are quite useful in providing humans with insights into the meaning of data and new useful inputs to supervised machine learning algorithms. As a kind of learning, it resembles the methods humans use to figure out that certain objects or events are from the same class, such as by observing the degree of similarity between objects. Some recommendation systems that you find on the web in the form of marketing automation are based on this type of learning.
- 3. Reinforcement learning:** When you present the algorithm with examples that lack labels, as in unsupervised learning. However, you can accompany an example with

positive or negative feedback according to the solution the algorithm proposes comes under the category of Reinforcement learning, which is connected to applications for which the algorithm must make decisions (so the product is prescriptive, not just descriptive, as in unsupervised learning), and the decisions bear consequences. In the human world, it is just like learning by trial and error. Errors help you learn because they have a penalty added (cost, loss of time, regret, pain, and so on), teaching you that a certain course of action is less likely to succeed than others. An interesting example of reinforcement learning occurs when computers learn to play video games by themselves.

In this case, an application presents the algorithm with examples of specific situations, such as having the gamer stuck in a maze while avoiding an enemy. The application lets the algorithm know the outcome of actions it takes, and learning occurs while trying to avoid what it discovers to be dangerous and to pursue survival. You can have a look at how the company Google DeepMind has created a reinforcement learning program that plays old Atari's video games. When watching the video, notice how the program is initially clumsy and unskilled but steadily improves with training until it becomes a champion.

1. Semi-supervised learning: where an incomplete training signal is given: a training set with some (often many) of the target outputs missing. There is a special case of this principle known as Transduction where the entire set of problem instances is known at learning time, except that part of the targets are missing.

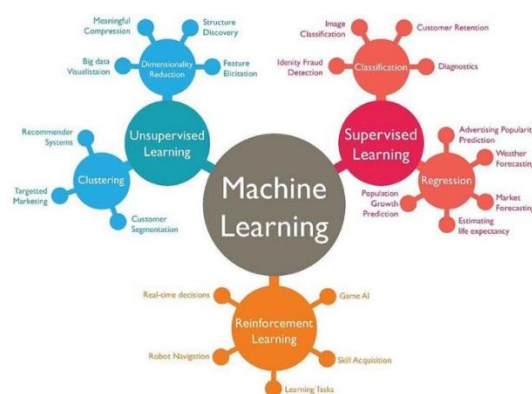


Figure 2: Classification of Machine learning.

Reinforcement learning

Reinforcement learning could be a machine learning training method supported rewarding desired behaviors and/or punishing undesired ones. In general, a reinforcement learning agent

is in a position to perceive and interpret its environment, take actions and learn through trial and error.

How does reinforcement learning work?

In reinforcement learning, developers devise a way of rewarding desired behaviors and punishing negative behaviors. This method assigns positive values to the required actions to encourage the agent and negative values to undesired behaviors. This programs the agent to hunt long-term and maximum overall reward to attain an optimal solution. These long-term goals help prevent the agent from stalling on lesser goals. With time, the agent learns to avoid the negative and seek the positive. This learning method has been adopted in computing (AI) as the way of directing unsupervised machine learning through rewards and penalties.

Applications and samples of reinforcement learning

While reinforcement learning has been a subject of much interest within the field of AI, its widespread, real-world adoption and application remain limited. Noting this, however, research papers abound on theoretical applications, and there are some successful use cases. Current use cases include, but don't seem to be limited to, the following:

- Gaming
- Resource management
- Personalized recommendations
- Robotics

Gaming is probably going the foremost common usage field for reinforcement learning. it's capable of achieving superhuman performance in numerous games. a standard example involves the sport Pac-Man.

A learning algorithm playing Pac-Man may need the power to maneuver in one in every of four possible directions, barring obstruction. From pixel data, an agent may well be given a numeric reward for the results of a unit of travel: 0 for empty space, 1 for pellets, 2 for fruit, 3 for power pellets, 4 for ghost post-power pellets, 5 for collecting all pellets and completing A level, and a 5-point deduction for collision with a ghost. The agent starts from randomized play and moves to more sophisticated play, learning the goal of getting all pellets to complete the extent. Given time, an agent might even learn tactics like conserving power pellets until needed for self-defense.

Reinforcement learning can operate in a very situation as long as a transparent reward will be applied. In enterprise resource management (ERM), reinforcement learning algorithms can allocate limited resources to different tasks as long as there's an overall goal it's trying to attain. A goal during this circumstance would be to avoid wasting time or conserve resources. In robotics, reinforcement learning has found its way into limited tests. this sort of machine learning can provide robots with the power to find out tasks an individual's teacher cannot demonstrate, to adapt a learned skill to a replacement task or to realize optimization despite a scarcity of analytic formulation available.

Reinforcement learning is additionally utilized in research, scientific theory, scientific theory, control theory, simulation-based optimization, multiagent systems, swarm intelligence, statistics and genetic algorithms.

Challenges of applying reinforcement learning

Reinforcement learning, while high in potential, is difficult to deploy and remains limited in its application. one in every of the barriers for deployment of this kind of machine learning is its reliance on exploration of the environment. For example, if you were to deploy a robot that was reliant on reinforcement learning to navigate a fancy physical environment, it'll seek new states and take different actions because it moves. it's difficult to consistently take the most effective actions during a real-world environment, however, thanks to how frequently the environment changes.

The time required to confirm the educational is finished properly through this method can limit its usefulness and be intensive on computing resources. because the training environment grows more complex, so too do demands on time and compute resources.

Supervised learning can deliver faster, more efficient results than reinforcement learning to companies if the right amount of knowledge is offered, because it will be employed with fewer resources.

Common reinforcement learning algorithms

Rather than pertaining to a selected algorithm, the sphere of reinforcement learning is formed from several algorithms that take somewhat different approaches. The differences are mainly because of their strategies for exploring their environments.

- State-action-reward-state-action (SARSA). This reinforcement learning algorithm starts by giving the agent what's referred to as a policy. The policy is basically a probability that tells it the percentages of certain actions leading to rewards, or beneficial states.
- Q-learning. This approach to reinforcement learning takes the alternative approach. The agent receives no policy, meaning its exploration of its environment is more self-directed.
- Deep Q-Networks. These algorithms utilize neural networks additionally to reinforcement learning techniques. They utilize the self-directed environment exploration of reinforcement learning. Future actions are supported a random sample of past beneficial actions learned by the neural network.

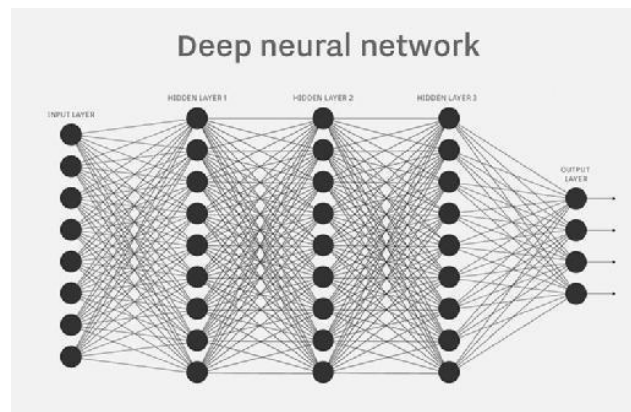


Figure 3: Deep Neural Network.

Reinforcement learning is taken into account its own branch of machine learning, though it does have some similarities to other forms of machine learning, which break down into the subsequent four domains: A reinforcement learning algorithm, or agent, learns by interacting with its environment. The agent receives rewards by performing correctly and penalties for performing incorrectly. The agent learns without intervention from a person's by maximizing its reward and minimizing its penalty. it's a sort of dynamic programming that trains algorithms employing a system of reward and punishment.

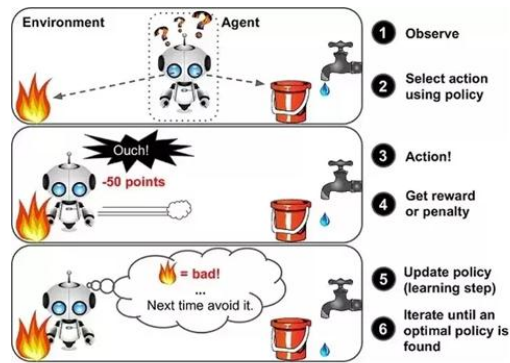


Figure 4: Example of Reinforcement Learning.

In the above example, we are able to see that the agent is given 2 options i.e. a path with water or a path with fire. A reinforcement algorithm works on reward a system i.e. if the agent uses the hearth path then the rewards are subtracted and agent tries to be told that it should avoid the fireplace path. If it had chosen the water path or the safe path then some points would be added to the reward points, the agent then would try and learn what path is safe and what path isn't. It's basically leveraging the rewards obtained, the agent improves its environment knowledge to pick the following action.

CONCLUSION

Reinforcement Learning addresses the matter of learning control strategies for autonomous agents with least or no data. RL algorithms are powerful in machine learning as collecting and labelling an outsized set of sample patterns cost quite data itself. The trial-and-error method because it attempts its task, with the goal of maximizing long-term reward can show better results here. Reinforcement learning is closely associated with dynamic programming approaches to Markov decision processes (MDP).

REFERENCES

1. Yu-Wei, Chiu (David Chiu) - Machine Learning With R Cookbook-Packt Publishing, 2015.
2. Sebastian Raschka, Vahid Mirjalili Python Machine Learning. 2nd ed-Packt Publishing, 2017.
3. Miroslav Kubat An Introduction to Machine Learning, 2018.
4. Drew Conway, John Myles White Machine Learning for Hackers-O'Reilly Media, 2012.
5. Ian H. Witten, Eibe Frank, Mark A. Hall Data Mining_ Practical Machine Learning Tools and Techniques, Third Edition-Morgan Kaufmann, 2011.