Available online www.ejaet.com

European Journal of Advances in Engineering and Technology, 2021, 8(6):35-41

Research Article ISSN: 2394 - 658X

The amazing realm of Machine Learning and its diverse array of types and categories

Ravi Shankar Koppula

*Ravikoppula100@gmail.com

ABSTRACT

This research paper explores the transformative impact of machine learning (ML) on various industries, emphasizing its crucial role in the development of artificial intelligence (AI) and data exploration. The increasing demand for advanced data processing methods necessitates the importance of AI and ML. ML, a solution that enables computers to learn from data without explicit programming, can be classified into three main types - supervised learning, unsupervised learning, and reinforcement learning. These types have distinct applications and significance in fields such as fraud detection, medical diagnostics, and web mining. Supervised learning is employed when labeled data is available, while unsupervised learning focuses on identifying patterns in unlabeled data and is often used for customer segmentation in marketing. Reinforcement learning, on the other hand, involves an agent learning to optimize long-term rewards through trial and error, making it valuable in dynamic and decision-making environments. This paper delves into various techniques, including K-means clustering, hierarchical clustering, Q-learning, and A3C, and elucidates their roles within these learning paradigms. The significance of evaluating models and efficient data processing is underscored in order to fully harness the potential of ML. Concrete examples in sectors such as aviation, health, and internet usage demonstrate how ML can reshape business models and societal norms. The paper presents comprehensive definitions of ML and its application areas and provides a comparison of supervised, unsupervised, and reinforcement learning, illustrating their unique contributions to both academia and industry. The insights gained from this study shed light on the notion that ML is not merely a theoretical concept but rather a practical tool with profound implications for marketing research and beyond.

Key words: Machine Learning, types of machine learning, Supervised Learning, Unsupervised Learning, Reinforcement Learning

INTRODUCTION

Artificial Intelligence (AI) is a domain of study in computer science that strives to cultivate computer systems capable of executing tasks that typically necessitate human intelligence. AI technology is fundamentally grounded in machine learning, where data is amassed from machines to endow them with intelligence. The crux of empowering AI systems with intelligence lies in engendering a supple logic that can adapt predicated on the outcomes of data processing. Once intelligence is attained, the subsequent phase of erecting Artificial Intelligence involves fostering erudition, adaptation, and self-organization within the machine. Over the preceding five decades, the focal point of AI technology has metamorphosed from proficiently comprehending voluminous amounts of information to encompassing domains such as heuristic search, problem-solving, reasoning, knowledge representation, planning, and machine learning. This alteration has considerably augmented the capabilities and potential applications of AI. One of the principal catalysts behind the precipitous expansion of AI technology is the quandary associated with the cost of item inspection. Myriad manufacturers aspire to automate processes; however, the exorbitant costs linked to automated inspection technology have hindered the widespread adoption of machine vision-based inspection systems. Notwithstanding, the development of intelligent applications and services utilizing AI has engendered momentous progress in the utilization of machine vision technology. This progress can be attained by amalgamating machine learning with evidence-based rules and preexisting knowledge to establish a pattern-based "intelligent" algorithm for automating item inspection. In our project, the primary emphasis will be on conducting experiments with various types of machine learning, including inductive learning through rule extraction, neural networks, and genetic algorithms. Additionally, we will appraise and compare the merits and demerits of these algorithms in diverse industries. Machine learning is a manifestation of artificial intelligence that empowers computers to learn sans explicit programming. Its focal point lies in fashioning computer programs that can adapt and mutate when exposed to new data. The process entails training a model, which is a system that formulates predictions predicated on the available data. The term "machine learning" was coined by Arthur Samuel in 1959, who delineated it as the study of conferring computers with the capability to learn autonomously. By way of illustration, in conventional software development, a program is furnished with a fixed set of data and executes a specific task predicated on said input. If the program generates erroneous output, the code is scrutinized, and the program is rerun with the same data. However, with machine learning, the software can self-correct, culminating in improved solutions. The input data is fed into the model, and both the program itself and its outputs are juxtaposed to a "ground truth." This evaluation is pivotal for testing the model's outputs and discerning areas where it can be enhanced through active or passive learning. In the present data-driven era, the significance of machine learning cannot be overstated. The availability of voluminous amounts of data from myriad sources has fueled the development of transformative technologies that can revolutionize industries, business models, and society as an entirety. The escalating intricacy of models and systems has engendered the need for expeditious and more efficacious methods of processing and analyzing data. Furthermore, advancements in computing power have enabled the automation of innumerable processes and the storage of data in previously unfathomable ways. Machine learning plays a pivotal role in facilitating these advancements, sanctioning novel perspectives in data exploration and the unearthing of latent potentials [1].

Definition

In this section, we will explain the concept itself by providing a clear definition of machine learning. We will also discuss how systems that utilize machine learning differ from those using traditional programming methods. Additionally, we will introduce various application areas where machine learning is applied. Machine learning is a technique within the realm of artificial intelligence (AI) that enables computers to learn without explicit programming. It is a captivating field closely linked with computer science and information technology. Moreover, machine learning profoundly impacts multiple aspects of our daily lives, spanning from business and science to health and security. It is worth emphasizing that machine learning operates based on models of the world, similar to how humans gain understanding from past experiences. This idea is crucial when defining machine learning. According to Mitchell's (1997) description, machine learning can be summarized as follows: "A computer program is deemed to learn from experience E in relation to a specific class of tasks T and evaluation measure P if its performance in tasks within T, as measured by P, improves with the acquisition of experience E." In this definition, E represents 'experience', P denotes 'performance', and T signifies 'task'.

Machine learning involves the process of identifying and analyzing patterns in datasets without the need for explicitly programmed instructions. This approach focuses on deriving functions that efficiently establish a connection between input and output data. The ultimate goal is to create a model or approximation of their relationship and connection, while also considering the inclusion of various attributes. These attributes are combined with input data, and algorithms are used to facilitate machine learning and produce an output, typically in the form of a model.[2]

Importance

The application of machine learning is extensive and expanding, and its significance is steadily growing. Undoubtedly, ML is a crucial technology in today's world. Its primary importance lies in its potential for career advancement. When data is inputted into a computer, an algorithm processes and analyzes the information, ultimately producing an output. This process is referred to as machine learning. Its significance extends to both research and industry fields. In research, machine learning aids in making advancements and accelerating progress. Additionally, speed plays a major role. For instance, if research solely relied on human intervention, it would consume considerable time to solve a single problem. Conversely, ML simplifies and expedites the process significantly, as a machine can think thousands and even millions of times faster than an average person.

The applications of machine learning are universal, from disease identification, such as cancer, to enhanced software verification. In the medical field, machine learning improves diagnostic accuracy, ensuring that serious conditions are not missed. Moreover, the implementation of ML also plays a crucial role in fraud detection. Banks, for instance, utilize ML to detect fraudulent credit card transactions, scanning billions of transactions annually. Cyber-attacks can also be identified using machine learning, as it monitors data transfer between computers. In the industrial sector, machine learning enables predictive maintenance, a term that encompasses the essence of this technology. Instead of repairing machines after they fail, predictive maintenance employs technological progress to repair them when necessary. Modern machinery can exhibit performance indicators, allowing for the prediction of potential issues. Consequently, many companies now rely on analog records to preemptively repair machines, rather than waiting for them to malfunction. Predictive maintenance, powered by machine learning, is becoming increasingly popular among large companies.

In the recent years, machine learning heavily transformed the computer science. It gave us new ways to use data, produce data and the most importantly, learn from data. This interesting capability is the reason why machine learning is applied in a numerous number of fields, from financial industry to medicine. In the modern world, where everything is driven by data like never before, this is the key and also the most important reason why machine learning is so important today. Also, with the help of machine learning, our intelligent systems will not only be able to process and analyze data, but they will also be able to actually "understand" it, in a very human sense of the word. With this capability, they can become an essential partner in decision making by providing useful information. For example, people are using machine learning to build intelligent "medical personal assistants" that will try to find the best ways to help doctors. By analyzing patient's data and the latest research results, these systems can provide suggestions to the doctors, such as what could be the most effective treatments for a certain illness. By actually applying machine learning algorithms to the system, it will be possible for it to learn from both the input data and the provided output and in turn, increase the quality of the suggestions it gives.

TYPES OF MACHINE LEARNING

Machine learning is generally split into three main categories based on the nature of the learning "signal" or "feedback" available to a learning system. These are Supervised Learning, Unsupervised Learning and Reinforcement Learning. In supervised learning, we have a dataset consisting of both the input features and the target. The target is something we want to predict and the input features are the properties which might help us predict the target. The supervised learning algorithm analyses the training data and produces an inferred function, which can be used for mapping new examples. The main aim is to construct a function which can accurately make predictions about the output values. In contrast, unsupervised learning is a type of machine learning algorithm used to draw inferences from datasets consisting of input data without labelled responses. The most common unsupervised learning method is cluster analysis, which is used for exploratory data analysis to find hidden patterns or grouping in data. The algorithm will use the input data to create a simplified model with the aim of classifying the different groupings or data sets within that model. Lastly, in reinforcement learning a computer employs trial and error to come up with a solution to a problem. The most important point in reinforcement learning is that it uses active exploration to obtain new data. This enables the algorithm to discover the best possible way to solve the problem. This is common in robotics, where the robot can learn without any human intervention. Every machine learning algorithm goes through these stages, though not all machine learning algorithms will take advantage of all three of them. However, the aforementioned algorithms are the most commonly used.

Machine Learning algorithms are often categorized as supervised or unsupervised. In a supervised learning model, the algorithm learns on a labeled dataset, that is, a dataset which known results are already tagged. This can be compared to unsupervised learning, where a model is given a dataset without giving it prior instructions, and the model then performs a sort of self-organization by adjusting itself to represent the data. Unsupervised learning is further divided into two categories: in clustering, the algorithm groups together similar types of data points based on specific features, whilst in dimension reduction, the algorithm discovers the underlying attributes or features that the data are represented by. Reinforcement learning is the third type of machine learning that involves a software agent that learns how to behave in an environment by performing actions and seeing the results, in order to achieve maximum cumulative reward. This type of learning is about making a sequence of decisions. It is similar to a supervised learning algorithm, but the reinforcement algorithm actually knows the correct output because it is getting rewards or penalties for the actions. It is learning from the training data that the supervised algorithm would use, because the algorithm has not been told what action to take, but it is making the decisions to receive the best reward possible, such as in a reinforcement simulation.[3]

Supervised Learning

Under the tutelage of a machine learning model, the supervised learning algorithm iteratively makes predictions on the training data and is relentlessly corrected by the teacher. Learning from the historical outputs, the model finds the optimal path to reduce loss function and generate insights by adjusting the weights that are multiplied to the different inputs. This type of learning method is fundamentally used in various applications such as fraud detection, image recognition or email filtering. Under supervised learning, there are two subcategories: 'classification' and 'regression'. In 'classification', the output variable is category type whilst in 'regression', the output is numerical - it allows the model to predict the numeric outcome by finding the best line that fits the value. This method will eventually break the dataset into smaller and smaller subsets whilst at the same time each decision tree generated will narrow down onto a specific value. With the input of historical and real-time flight route and its speed and location, supervised learning can be deployed in aviation such as to predict plane faults before it happens. The output of supervised learning in this case will typically be a yes/no answer, and the model might look at many features that characterize inbound emails such as the sender, the email's route and email content itself. The model can do this because of prior teaching that spam emails usually contain certain traits. By providing the model with a series of historical data and the results of diagnoses, supervised learning could be used to identify and customize potential treatment approaches for patients.

The training data that has been used is providing the correct answers, which they refer to as labels. Through this data, the learning algorithm can make predictions on a new set of data. This data has generally been used in classification and regression. In classification, the output variable is a category, for example showing us either email is spam or not spam. This means that we are likely to use classification learning where the data we have is used to categorise something, mostly in groups of large data set obtained. On the other hand, Supervised Learning is used when the output is a real value, such as "weight". This type of learning has been used in regression because it has been categorical. An example of supervised learning used in medical diagnostics is to allow doctors to provide the most effective treatment. By using supervised learning, they are taught how the functions of treatment in a wide range of patients.[4]

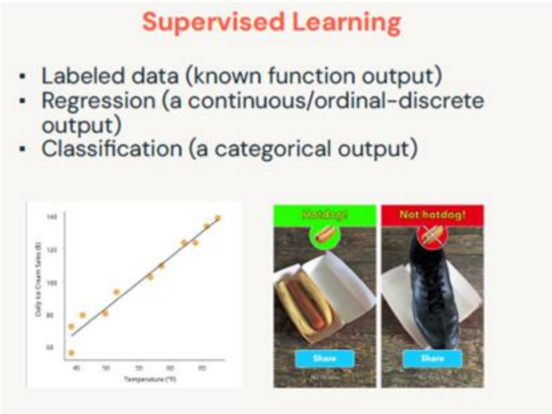

Unsupervised Learning

Unsupervised learning is the training of an artificial intelligence (AI) algorithm using information that is neither classified nor labeled and allowing the algorithm to act on that information without guidance. Here the task of the machine is to group unsorted information according to similarities, patterns and differences without any prior training of data. The main goal of unsupervised learning is to explore the underlying structure of the data. This can help the data scientist to understand the properties of the data. This technique is usually used in complex processing tasks such as – when the data is too complex for labeled data, large scale data or for a higher-level feature learning.

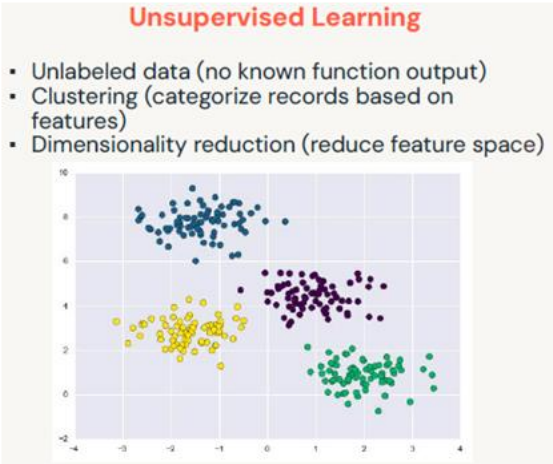

There are several unsupervised learning methods. Some of them are mainly used for clustering and other is used for dimensionality reduction. The more commonly used unsupervised methods are - 1) Clustering methods and 2) Principal Components Analysis (PCA) And these two are mainly used for finding patterns of data and reduce the dimension of data. The below table describes some of the applications of the unsupervised learning techniques. And these applications have used the technique to identify patterns and give insight about the data. Applications Description Web mining Web mining is the application of unsupervised learning techniques for finding patterns in large data sets that can be used to understand and - Extract information from web. Identify similar contents from different websites. Marketing Research Marketing Research is the process to analyze the consumer behavior and - Have used the technique to search for certain segments of similarity between different customer population.

In unsupervised learning, on the other hand, the data provided is neither labeled nor classified. This means that the algorithm can act on the data without any guidance; essentially, the algorithm attempts to learn the patterns and the features in the input. Unsupervised learning is used for clustering - that is for finding groups within the data. Unsupervised learning is also used for dimensionality reduction. This is useful because it often makes the data easier to be visualized. Think about it, most of us can only visualize in 3 dimensions, yet the data provided can often be lots of dimensions - each separate variable forms a new dimension, so it can become quite complex. If you can apply dimensionality reduction so a complex data set can be understood in 3 dimensions or 2, then human understanding and visualization of the data becomes much more feasible and clearer. There are many algorithms for unsupervised learning and the complexity and difficulty can vary significantly. One of the most common forms of unsupervised learning is called K-means clustering. This is where the algorithm tries to split the data into K different clusters, where each point in the cluster is closer to its mean than a mean in any other cluster. The ‘K’ is essentially the user proving to the algorithm how many clusters are believed to be within the data. The algorithm then tries to best fit these clusters to the data. The algorithm does this by iteratively reassigning points in the clusters and recalculating the mean, until the variation in the assigned clusters is minimized. This is a hard form of clustering, which means that each data point is assigned to only one cluster and this type of algorithm is often used in customer segmentation, very useful in marketing analysis. However, other types if unsupervised learning exist - for example hierarchical clustering. This is similar to K-means, but instead of being pre-determined how many clusters will be found in the data, the algorithm successively groups the points into a tree of clusters - it’s a bit like a family tree. This is called hierarchical because you have a hierarchy of clusters - you can have a few large clusters, with small ones inside, and then each small cluster could even be broken into smaller clusters. This is known as a soft clustering technique, where in the end the data points can end up belonging to more than one cluster - it potentially might have better applicability to various data sets or it may be easier to find natural clusters using the visualization the hierarchy produces. [5]

Reinforcement Learning

Reinforcement Learning (RL) is a segment of machine learning. It sets up a representation of an intuitive agent, for example, a robot, who adapts so as to accomplish a perfect reward. At every time, the agent plays out an activity contingent upon its present state, after which it advances to a resulting state and gets a similar reward.

This is a specific process for machine learning calculations where the machine is trained to settle on ideal choices to arrive at a specific objective. The operator learns by methods for experimentation, gaining intense lessons on that way and activities prompt the best reward. Normally, both the reward and the penalty are characterized by the human creator in the issue conditions, and the issue condition itself is driven by the machine control activities. It is imperative to take note of that, albeit after some time the algorithmic specialist will learn propensities that rise to the best reward, it doesn't predict or reallocate its gained information. The RL specialist goes through the world with the exclusive center placed on picking up the current best reward. This means the machine doesn't typecast or supernaturally anticipate what courses of action and activities lead to ideal outcomes; it learns them over rehashed tests in the issue space.

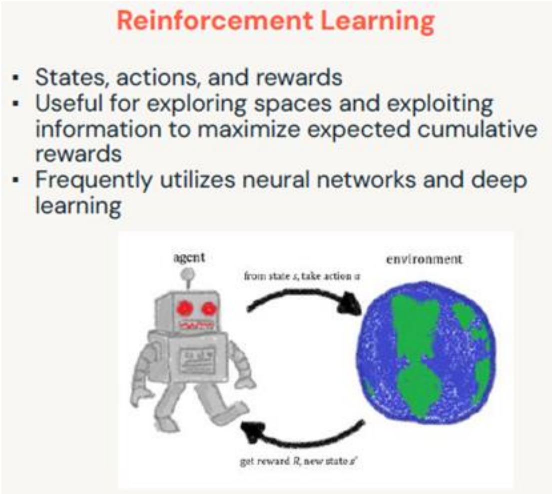

When communication is occurring between a few zones of the machine brain or between the machine mind and the outside world, these associations are dynamic however intrinsically constrained to the on-going operations of the machine. RL is totally different to regulated and unsupervised learning. With RL, the information focuses and the reward function are obscure to the agent. In supervised learning, the learning algorithm creates a model from information focal points on high-dimensional information that outlines to an optimal output. By correlating factors and yields and using repetitive re-weightings of the contribution of every factor, managed learning algorithms fill the gaps in learning to predict yields founded on sources of info. In unsupervised learning, meanwhile, no unique outcomes administer the propriety of the learned data examples. This is a basic contrast from bolstered learning. In supervised learning the operator realizes what activity ought to be taken, without a reward capacity code; the operator just processes the information focuses and their features to "learn a topology" that prompts the ideal work. With supervised learning, the objective is to restore the most helpful target work counterpart to each point in the information space. However, in RL, the information illustrate and the reward work all elements of the operator's extraordinary work.

Reinforcement Learning (RL) is a type of machine learning where an agent learns how to behave in an environment by performing certain actions and observing the rewards it gets from those actions. Starts with the agent in a state where it must decide what action to take from that state. Performing the action transitions the agent to a new state and it also receives a real value as a reward from performing that action. This allows the agent to learn from positive rewards, that encourage certain actions to be performed, and from negative rewards or penalties, which discourage certain actions. By accumulating these rewards over a sequence of state/action pairs, the agent can determine which sequence of actions leads to the highest expected long-term reward. This is referred to as a policy and the overarching goal of reinforcement learning is to find the optimal policy for the given environment. Optimal in this sense means that following that policy will lead the agent to achieve the greatest location of rewards over time, from that starting state. One way of doing this is by using what's called Q-learning. Q-Learning is a model-free and off-policy algorithm that seeks to find the best action to take given a particular state. It's a value-based method that aims to approximate the so-called action value function Q, which is a function of both the state and action, and represents the expected return of taking a given action from that state, and then following the optimal policy thereafter. By using Q-learning, the agent can learn an optimal policy without needing to learn the transition probabilities from each state, as would be required for a model-based method such as the probability-based family of solution methods called and A3C is also a model-free method in the sense that it doesn't rely on build and use a model of the environment in the traditional sense Japanese model-based methods do. Unlike Q-Learning, however, A3C is an on-policy algorithm, meaning that the policy being learned and the policy followed by the agent during action selection are the same. The key advantage of A3C over Q-learning is that it operates with a large degree of parallelism, and it's well-suited for environments with high-dimensional state spaces and a continuous action domain. [6]

SUMMARY

This paper takes a deep dive into the world of machine learning, highlighting its incredible potential to enable computers to learn on their own without needing explicit programming. It emphasizes the transformative impact that machine learning has had on various industries. The paper delves into different types of machine learning, including supervised, unsupervised, and reinforcement learning. Supervised learning is examined in detail, focusing on how it utilizes labeled datasets to train predictive models. On the other hand, unsupervised learning focuses on discovering hidden patterns in unlabeled data, utilizing techniques such as K-means and hierarchical clustering. This type of learning is particularly crucial in marketing, as it assists with customer segmentation and offers a more nuanced understanding of data relationships. Reinforcement learning, with its reward-based approach to learning, is explored through methods like Q-learning and A3C. The latter is especially suited for complex environments with high-dimensional state spaces. Reinforcement learning stands out for its unique approach to data usage and learning processes, setting it apart from the other two types of machine learning. Throughout the paper, the author emphasizes the pivotal role that machine learning plays in advancing research and decision-making. Numerous applications are highlighted, ranging from medical diagnostics to fraud detection and web mining. These examples illustrate how machine learning has become an indispensable tool across a wide range of fields, improving outcomes and enabling new discoveries.

REFERENCES

[1]. A. Dey, “Machine Learning Algorithms: A Review.” Available: https://ijcsit.com/docs/Volume%207/vol7issue3/ijcsit2016070332.pdf

[2]. L. D. Jones, D. Golan, S. A. Hanna, and M. Ramachandran, “Artificial intelligence, machine learning and the evolution of healthcare,” Bone & Joint Research, vol. 7, no. 3, pp. 223–225, May 2018, doi: https://doi.org/10.1302/2046-3758.73.BJR-2017-0147.R1.

[3]. P. Bradshaw, “Data journalism’s AI opportunity: the 3 different types of machine learning & how they have already been used,” Online Journalism Blog, Dec. 14, 2017. https://onlinejournalismblog.com/2017/12/14/data-journalisms-ai-opportunity-the-3-different-types-of-machine-learning-how-they-have-already-been-used

[4]. J E T Akinsola, “Supervised Machine Learning Algorithms: Classification and Comparison,” ResearchGate, Jun. 08, 2017. https://www.researchgate.net/publication/318338750_Supervised_Machine_Learning_Algorithms_Classification_and_Comparison

[5]. S. Jaeger, S. Fulle, and S. Turk, “Mol2vec: Unsupervised Machine Learning Approach with Chemical Intuition,” Journal of Chemical Information and Modeling, vol. 58, no. 1, pp. 27–35, Jan. 2018, doi: https://doi.org/10.1021/acs.jcim.7b00616.

[6]. P. Henderson, R. Islam, P. Bachman, J. Pineau, D. Precup, and D. Meger, “Deep Reinforcement Learning That Matters,” Proceedings of the AAAI Conference on Artificial Intelligence, vol. 32, no. 1, Apr. 2018, Available: https://ojs.aaai.org/index.php/AAAI/article/view/11694