JETIR.ORG ISSN: 2349-5162 | ESTD Year : 2014 | Monthly Issue
JOURNAL OF EMERGING TECHNOLOGIES AND INNOVATIVE RESEARCH (JETIR)
An International Scholarly Open Access, Peer-reviewed, Refereed Journal

Artificial Intelligence

AI Marvels: Exploring Convolutional Neural Networks, Deep Q-Learning, and A3C

[1]Saloni Kumari

[1]Software Engineer
[1]Technology and Consulting,
[1]EY (Ernst & Young), Hyderabad, India

Abstract: Artificial intelligence (AI) is the emulation of human intelligence in devices intended to act and think like people. It entails the creation of computer programs and algorithms that can carry out operations that ordinarily need human intellect, such as speech recognition, visual perception, decision-making, and language translation. This paper discusses the Artificial Neural Networks, Convolutional Neural Networks, Reinforcement Learning, Deep Q-learning Intuition, Deep Convolutional Q-learning Intuition and A3C Intuition.

IndexTerms **- ANN, CNN, Reinforcement Learning, Deep Q-learning, A3C Intuition**

I. INTRODUCTION

Artificial intelligence (AI) is the replication of human intelligence in machines that are created to think and learn in ways that are comparable to those of humans. Artificial intelligence (AI) systems can be created to carry out operations that ordinarily require human intelligence, such as comprehending spoken language, identifying objects in photographs, and making judgement calls. AI comes in a variety of forms, including:

Reactive machines: These AI systems lack the capacity to store memories or draw conclusions from the past to guide judgements made now.

Limited memory: These AI systems have a limited ability to construct memories, but they can use prior experiences to inform current judgements.

Theory of mind: These AI systems can comprehend human emotions, beliefs, and intentions.

Self-aware: These AI systems have self-awareness and are genuinely conscious.

AI can be used in a wide range of industries, including manufacturing, banking, transportation, and health care. AI research aims to develop machines that can carry out any intellectual work that a human is capable of.

There are numerous actual instances of AI being employed in numerous fields and applications. Several instances include:

Healthcare: To identify diseases and cancers, AI is being used to evaluate medical imagery like CT and MRI scans. In order to pinpoint risk variables and forecast results, it is also used to evaluate patient data from sources like electronic health records.

Finance: AI is used to spot fraudulent credit applications and transactions, as well as to analyze financial data, such as stock prices and trade patterns, to forecast the future and spot trends.

Retail: To personalize recommendations and enhance the customer experience, AI is being used to evaluate user data and behaviors, such as browsing and purchasing history.

Artificial intelligence is being utilized to steer autonomous automobiles, trucks, and drones. It makes it possible for vehicles to process sensor data from tools like cameras, lidar, radar, and GPS to comprehend their surroundings and make choices.

Manufacturing: AI is being used to monitor and operate industrial equipment, such as robots and conveyor belts, as well as to optimize production processes, such as scheduling and logistics.

Natural Language Processing: Businesses utilize chatbots and virtual assistants that are AI-powered to help customers with sales and customer support inquiries.

Speech recognition: To convert spoken words into text, virtual assistants like Amazon's Alexa and Google's Home, as well as mobile apps, use AI-powered speech recognition.

These are only a few real-world applications of AI, although there are many more. In many fields and applications, AI is becoming more and more common, and this trend is likely to continue in the future.

II. RESEARCH METHODOLOGY

2.1 Artificial Neural Networks

A sort of machine learning model called an artificial neural network (ANN) is motivated by the form and operation of biological neural networks. Layers of interconnected "neurons" make up ANNs, which process and send information. Each neuron gets information from other neurons, processes it, and then transmits the results to other neurons.

The artificial neuron, which performs a straightforward mathematical operation on the input it receives and generates an output, is the fundamental component of an ANN. The artificial neurons' input and output are numerical, and the action taken is referred to as an activation function. The network can learn non-linear correlations between inputs and outputs thanks to the activation function.

By changing the weights on the connections between neurons, a neural network can be trained to output the desired result for a given input. This is accomplished by a procedure known as backpropagation, which determines the ideal set of weights using the gradient descent technique.

ANNs come in a variety of forms, including convolutional, recurrent, and feedforward networks. These various architectures can handle a variety of tasks and data, including prediction, natural language processing, and picture recognition.

2.2 Convolutional Neural Networks

An artificial neural network called a convolutional neural network (CNN) is made to process and interpret visual data, including pictures and movies. Convolutional layers, which apply a series of filters to the input data and enable the network to extract relevant features like edges, textures, and forms, are the distinguishing characteristic of CNNs.

The input data, like an image, is passed through several layers of filters in a CNN. Each filter moves over the image, examining discrete areas to create a fresh feature map. Following that, these feature maps are subjected to several layers of pooling, which lowers the spatial resolution of the data and increases its resistance to translations and distortions.

The final classification or regression operation is carried out by fully connected layers after the feature maps have passed through the convolutional and pooling layers. Like conventional neural networks, these layers link every neuron in one layer to every neuron in the following layer.

Wide-ranging applications like image classification, object detection, semantic segmentation, and video analysis have all seen great success using CNNs. This is so that CNNs, which can achieve high accuracy even with a relatively limited quantity of training data, can extract relevant features from photos. Additionally, CNNs are an effective tool for computer vision problems because they can be taught to handle large-scale picture datasets.

2.3 Fundamentals of Reinforcement Learning

The goal of Reinforcement Learning (RL), a subset of machine learning, is to educate agents to choose actions that will maximize a cumulative reward. It is employed to resolve issues where an agent interacts with the environment and gradually learns to respond in ways that produce the greatest rewards.

The core ideas behind RL include:

1. Agent: The thing that acts and engages with its surroundings.
2. Environment: The system with which the agent interacts and is provided with feedback.
3. State: A depiction of the environment's current state or circumstances.
4. Action: The decision the agent takes in reaction to the situation as it is.
5. Reward: The input the agent receives from the environment in response to its behaviours.
6. Policy: The plan or approach an actor uses to decide which actions to do.
7. Value is the anticipated future gain from a particular state or action.

In RL, the agent learns through contact with the environment, making mistakes along the way and using the reward as feedback to alter its strategy and gradually enhance the cumulative reward. The agent's objective is to identify the best course of action that maximizes the projected total reward. By estimating the value of various actions and states and updating the policy accordingly, techniques like Q-learning and SARSA are used to do this.

The use of reinforcement learning is widespread in fields including robotics, gaming, and recommendation engines.

2.4 Deep Q-Learning Intuition

Finding the best action-selection strategy for an agent interacting with the environment is done using the Q-learning method, which is a form of reinforcement learning algorithm. Estimating the "quality" or "value" of each potential action at a given state and choosing the action that maximizes this value are the central idea of Q-learning. To keep track of the estimated values, the method uses a Q-table, which it updates over time as the agent gains more knowledge about the environment. The idea behind Q-learning is to do the action, update the Q-table based on the reward received, and then learn the optimum action to take in a specific situation. By continuing to learn in this manner, the agent eventually approaches the ideal course of action.

In contrast to the standard Q-learning method, deep Q-learning represents the Q-table using a deep neural network. The idea behind deep Q-learning is that by teaching the neural network a more intricate and precise representation of the state-action value function, the agent will be able to more closely approach the optimal course of action.

The neural network, also referred to as the Q-network, in Deep Q-learning receives the environment's current state and produces a set of estimated values for each potential action. The Q-network is then updated with the observed reward and the new state, and the agent chooses the action with the highest estimated value. Using this method frequently, the Q-network is trained and eventually learns to approximate the actual Q-values of the states and actions. High-dimensional and continuous state spaces, which are intractable with conventional table-based Q-learning, can be handled using this method.

2.5 Deep Convolutional Q-Learning Intuition

Convolutional neural networks (CNNs) are used in Deep Convolutional Q-learning (DCQ), a type of deep Q-learning that processes visual input from the environment. According to the logic underpinning DCQ, CNNs are especially well-suited for processing picture data since they can pick up helpful information from the input that may be used to estimate how valuable an action will be.

To handle the unprocessed visual input from the environment, such as a video frame or screenshot, DCQ uses the CNN. Edges, textures, and other important features are extracted from the input by the CNN and passed to the Q-network. The agent chooses the action with the highest estimated value after the Q-network analyses these features to determine the value of each potential action. The CNN is trained concurrently with the Q-network and updated with the observed reward and the new state.

The agent's capacity to approximatively determine the best course of action can be enhanced by DCQ since it enables the agent to process high-dimensional visual input and learn useful aspects from the picture data. It is frequently used to solve visual input-related issues in robotic manipulation, video games, and autonomous cars..

2.6 A3C Intuition

A3C (Asynchronous Advantage Actor-Critic) is a variation of the actor-critic algorithm, a kind of reinforcement learning system that employs two different neural networks: one to pick actions and the other to judge the worth of those actions (the critic) (the actor). The idea behind A3C is to simultaneously and concurrently train the actor and critic while utilizing numerous instances, or "workers," of the environment. To gather data and update the shared actor and critic networks, A3C uses several worker threads that operate independently and asynchronously. Each worker keeps its own set of actor network parameters and interacts with its own replica of the environment. The gradients of the critic network, which is updated using the experiences of all the workers, are used to update the actor network of the worker.

The primary idea behind A3C is that it can considerably speed up the learning process and converge to an optimal policy much faster by having many workers working in parallel. Additionally, it can lessen variance and increase stability of the learning process by independently and concurrently updating the actor and critic networks.

When working with situations that have high-dimensional, continuous state spaces that take a long time to investigate, A3C is especially helpful. Applications like robots, video games, and simulations frequently employ A3C.

2.7 Comparison of the Algorithms

ANNs (Artificial Neural Networks):

1. ANNs are the core of deep learning and are supposed to imitate the neuronal structure of the human brain.
2. They are made up of interconnected layers of neurons that process input data and predict output.
3. They are appropriate for a variety of applications, including classification, regression, and pattern recognition.
4. Backpropagation and optimization methods are commonly used in training ANNs to alter the model's weights and biases.

CNNs (Convolutional Neural Networks):

1. CNNs are a sort of neural network that is specifically built for image recognition and computer vision tasks.
2. Convolutional layers are used to discover patterns and characteristics in photos automatically.
3. CNNs are extremely successful because they can capture local spatial dependencies using convolution processes.
4. They have demonstrated cutting-edge performance in tasks such as picture classification, object detection, and image segmentation.

RL (Reinforcement Learning):

1. RL is a subfield of machine learning in which an agent learns to operate in a given environment in order to maximize cumulative rewards.
2. The agent interacts with its surroundings and receives feedback in the form of rewards or punishments.
3. The goal of RL is to find the best policy to guide the agent's decision-making process.
4. RL has been effectively used to activities including as gaming, robotics, and self-driving vehicles.

Deep Q-Learning:

1. Deep Q-Learning is a sort of reinforcement learning algorithm that blends Q-Learning and deep neural networks.
2. It estimates the Q-values for various actions in each state using deep neural networks (typically CNNs) as function approximators.
3. DQN solves some of traditional Q-Learning's shortcomings in dealing with high-dimensional state spaces.
4. It has demonstrated extraordinary success in playing Atari games and other difficult activitie

Deep Convolutional Q-Learning:

1. Deep Convolutional Q-Learning is a DQN variation optimized for visual inputs like images.
2. It processes the visual input and estimates Q-values for various actions using CNN architectures as function approximators.
3. This method has contributed to superhuman performance in a variety of video games.

A3C Intuition (Asynchronous Advantage Actor-Critical):

1. A3C is an advanced RL method that trains using the actor-critic architecture and asynchronous parallelism.
2. It employs many agents that study different sections of the environment at the same time and update a central global model.
3. The actor network chooses actions, while the critic network estimates the value function to evaluate the chosen actions.
4. A3C has shown faster convergence and better stability than other RL approaches.

III. RESULTS AND DISCUSSION

Deep learning is built on Artificial neuronal Networks (ANNs), which simulate the neuronal structure of the human brain. They are made up of interconnected layers of neurons that excel at tasks such as classification and regression. Convolutional Neural Networks (CNNs) are computer vision specialized networks that use convolutional layers to recognize patterns in images, producing state-of-the-art performance in applications such as image identification. Reinforcement Learning (RL) is the process through which an agent learns to make decisions in an environment to maximize rewards. Deep Q-Learning blends RL with deep neural networks, resulting in breakthroughs in games such as Atari. Deep Convolutional Q-Learning improves DQNs by employing CNNs for visual input, driving video game performance to new heights. A3C is a sophisticated RL algorithm that uses parallelism and the actor-critic architecture to achieve faster convergence and increased stability. The algorithm chosen is determined on the problem and data properties.

IV. ACKNOWLEDGMENT

I would like to express our sincere gratitude to my organization EY (Ernst & Young) for unwavering guidance, invaluable insights, and constant encouragement throughout this journey. Their valuable input and feedback significantly improved the quality of this research. I would like to acknowledge the contributions of my research colleagues and friends who provided valuable feedback, engaging discussions, and constructive criticism. Their diverse perspectives enriched this work significantly.

REFERENCES

[1] "Understanding Emerging Technology: Artificial Intelligence." CompTIA. November 2017.
[2] "AI Is Continuing Its Assault on Radiologists." MIT Technology Review. January 2018.
[3] "How artificial intelligence will affect the future of energy and climate." Brookings Institution. January 2019.
[4] "Forget algorithms. The future of AI is hardware!" HuffPo. January 2018.
[5] "There's Just No Doubt That It Will Change the World': David Chalmers on V.R. and A.I." The New York Times. June 2019
[6] "Some studies in machine learning using the game of checkers." A. L. Samuel in IBM Journal of Research and Development. 1959.