JYOTI NIVAS COLLEGE
BENGALURU,INDIA
POST GRADUTE CENTER

DEPARTMENT OF MCA

TECH-ON-TAP

E-JOURNAL

ON

ARTIFICIAL INTELLIGENCE

ISSUE 3

AUGUST 2023

INDEX

Information security

Lakshmi Lavanya
Yuktha

The definition of Information Security states Preservation of confidentiality, integrity and availability of information. It is part of information risk management. It typically involves preventing or reducing the probability of unauthorized or inappropriate access to data or the unlawful use, disclosure, disruption, deletion, corruption, modification, inspection, recording, or devaluation of information. It also involves actions intended to reduce the adverse impacts of such incidents.

OVERVIEW:

At the core of information security is information assurance, the act of maintaining the confidentiality, integrity, and availability (CIA) of information, ensuring that information is not compromised in any way when critical issues arise.

These issues include but are not limited to natural disasters, computer/server malfunction, and physical theft. While paper-based business operations are still prevalent, requiring their own set of information security practices, enterprise digital initiatives are increasingly being emphasized, with information assurance now typically being dealt with by information technology (IT) security specialists. These specialists apply information security to technology (most often some form of computer system).

Information Security Attributes:

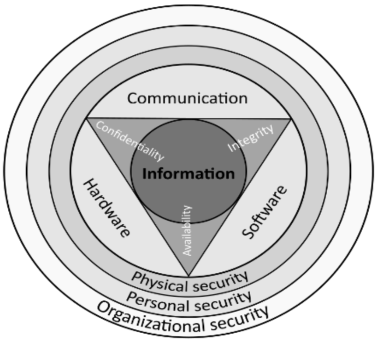

Confidentiality, Integrity and Availability (CIA). Information Systems are composed in three main portions, hardware, software and communications with the purpose to help identify and apply information security industry standards, as mechanisms of protection and prevention, at three levels or layers: physical, personal and organizational. Essentially, procedures or policies are implemented to tell administrators, users and operators how to use products to ensure information security within the organizations.

Possible responses to a security threat or risk are:

- reduce/mitigate – implement safeguards and countermeasures to eliminate vulnerabilities or block threats
- assign/transfer – place the cost of the threat onto another entity or organization such as purchasing insurance or outsourcing
- accept – evaluate if the cost of the countermeasure outweighs the possible cost of loss due to the threat

Risk management:

Risk is the likelihood that something bad will happen that causes harm to an informational asset (or the loss of the asset). A vulnerability is a weakness that could be used to endanger or cause harm to an informational asset.

The Certified Information Systems Auditor (CISA) Review Manual 2006 defines risk management as "the process of identifying vulnerabilities and threats to the information resources used by an organization in achieving business objectives, and deciding what countermeasures, if any, to take in reducing risk to an acceptable level, based on the value of the information resource to the organization."

Cryptography:

Information security uses cryptography to transform usable information into a form that renders it unusable by anyone other than an authorized user; this process is called encryption.

Security governance:

- An enterprise-wide issue
- Leaders are accountable
- Viewed as a business requirement
- Risk-based
- Roles, responsibilities, and segregation of duties defined
- Addressed and enforced in policy
- Adequate resources committed
- Staff aware and trained
- A development life cycle requirement
- Planned, managed, measurable, and measured
- Reviewed and audited

ROBOTICS

Bhumica R
Kavya L

What is robotics?
Robotics is the branch of engineering and computer sciences where machines are built to perform programmed tasks without further human intervention.

What is the role of artificial intelligence in robots?
Despite this, robotics and artificial intelligence can coexist.
Projects using AI in robotics are in the minority, but such designs are likely to become more common in future as our AI systems become more sophisticated. Here are some examples of existing robots that use AI.

EXAMPLES:
Agriculture: weed control, cloud seeding, planting seeds, harvesting, environmental monitoring and soil analysis.
Manufacturing: material handling, processing operations and assembly and inspection.
Healthcare: Increased efficiency of the diagnostic process, Reduced overall costs of running the business, safer surgeries.
Business

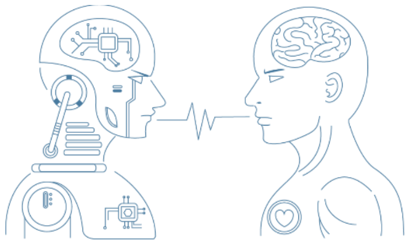

Are robotics and artificial intelligence the same thing?
Robotics and artificial intelligence are two different things.
Artificial intelligence is where systems emulate the human mind to learn, solve problems and make decisions on the fly, without needing the instructions specifically programmed.
Robotics is where robots are built and programmed to perform very specific duties.
In most cases, this simply doesn't require artificial intelligence, as the tasks performed are predictable, repetitive and don't need additional 'thought'.

Future of Robotics and Artificial Intelligence
Robotics and artificial intelligence are two related but entirely different fields.
Robotics involves the creation of robots to perform tasks without further intervention, while AI is how systems emulate the human mind to make decisions and 'learn.'

Types of robots:
- Pre-programmed robots
- Humanoid robots
- Autonomous robots
- Tele operated robots
- Augmenting robots

ALGORITHMIC GAME THEORY AND COMPUTATIONAL MECHANISM DESIGN

Neha
Nayana

Algorithmic Game theory is about strategic interactions among intelligent individuals, and mechanism design is about creating effective incentives in economic settings. Together, they're fascinating ways to understand human behavior and the challenges of designing and building systems.

Algorithmic game theory (AGT) is a way of analyzing social interactions that use mathematical models to predict the strategies that individuals will adopt in any given situation. A simple game theory model can predict human behavior in many situations. But the surprising thing is that this same model can also explain the complex, self-organizing systems that power the World Wide Web.

Algorithmic Game Theory:

Algorithms are the basis for much of our digital world — everything from social media networks to search engines to online advertising to the world wide web itself. These algorithms run on computational mechanisms. The study of algorithmic game theory and computational mechanism design is a relatively new field within computer science, but it has already made some profound contributions to a wide variety of fields, including politics, economics, and warfare.

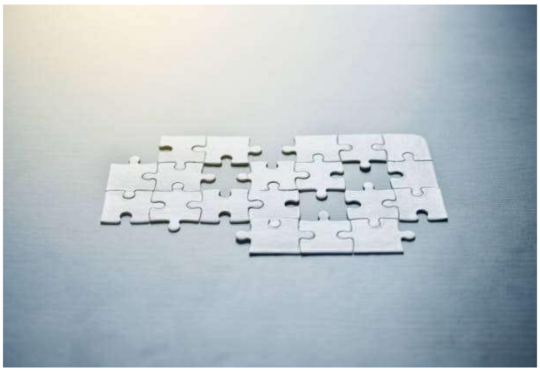

Algorithmic game theory is one of the most fascinating branches of computer science. It can be used to study games such as rock-paper-scissors, chess, and tic-tac-toe. And it has applications in engineering, economics, and psychology. Algorithmic Game Theory is a field in economics and computer science studying situations where agents interact with each other based on some shared objective.

Computational Mechanism Design:

A computational mechanism is an algorithm that is used to solve some specific problem. Computational mechanism design (CMD) is the act of designing algorithms that solve problems, specifically in the context of the Internet of Things (IoT).

CMD works by analyzing and understanding the behavior of these devices to create mechanisms that are capable of solving any particular problem that those devices may encounter. In order for a solution to be viable, there must be a clear problem, a good solution, and a market for the solution.

Why Are Algorithmic Game Theory and Computational Mechanism Design Related?

Game theory (GT) is the study of how people can make decisions that maximize their outcomes. Algorithmic game theory (AGT) is a branch of GT focused on how algorithms can be used to help people decide what to do or what strategy to adopt. Mechanism design (MD) is a subfield of AGT that focuses on designing mechanisms that implement desired outcomes. The goals of MD are to create mechanisms that are simple, predictable, and fair.

The key difference between game theory and computational mechanism design is that the former deals with human interactions. But the two are related because of the shared focus on studying the ways people interact with each other. So, in short, if we understand human behavior, we can better understand how a computer will behave, and vice versa.

Game Theory Applied to Real-World Problems

Game theory is the study of strategic interactions. When applied to the real world, the game theory looks at how individuals within a society make decisions in order to maximize their own success and minimize their own risk. The word “game” in game theory refers to a situation in which the participants are making decisions and there are consequences for their choices.

Game theory provides insight into why people behave in certain ways. A very famous game-theoretic concept is the prisoner’s dilemma. This concept states that there is an inherent conflict between two individuals’ interests. The individual in the scenario believes he or she can only benefit by cooperating. Still, that cooperation will not benefit the other individual who has the choice to cooperate or not. In a prisoner’s dilemma situation, a rational person would not choose to cooperate. Game theory helps us understand why people choose not to cooperate.

DECISION MANAGEMENT SYSTEM

PRATHYUSHA
JAHNAVI

Decision management is the process of designing, building, and managing automated decision-making systems. A business can make thousands of decisions a day; therefore, decisions are central to any business or organization. The challenge comes in when you have to take a significant number of factors into consideration at once that could influence a decision such as fluctuating prices, new products, current events, and more. That's where a decision management tool can greatly increase the efficiency of a company's decisions. A business can use decision management to help manage customer, supplier, and employee interactions. Using decision management, businesses change the way they approach their workflow. Incorporating and leveraging the power of big_data, decision management helps businesses meet operational requirements and user expectations.

Decision management finds its place within analytics—the discipline that uses business logic as well as data mathematics to help provide insights that make better decisions.

Analytics can be one of four types:

1. Descriptive: Tells you about something that has occurred, but not what you can do about it.
2. Diagnostic: Where you can explore the reasons why something occurred.
3. Predictive: This data models the likelihood that an event may happen in the future.
4. Prescriptive analytics: Combines all of the above but takes an extra step to make recommendations on how to respond or you can automate responses.

Decision management forms a part of prescriptive analytics. It helps organizations make operational decisions and trigger a response where needed as well. Decision management becomes a discipline which has sets of tools and techniques that enable businesses to make better decisions resulting in improved outcomes.

Areas of Decision Management

The goal of decision management is to enhance business operations intelligence by ensuring quick, consistent, and accurate fact-based decisions. The quality of structured operational decisions, no matter how complex, should be constantly improving. There are five areas that affect decision management:

- **Data and analytics**: Data is accessed and processed with the help of descriptive, diagnostic, and predictive techniques. You need strong data quality as a basis for accurate decision making, and the outcomes of those decisions affect the data as well.
- **Business process management**: Managing human tasks and the sequence of business process automation and task management. The information from staff helps to make better decisions, and their roles are enhanced as a result.
- **Operations research**: Optimizing and managing various goals based on standards and priorities that can be modeled. Decision management analyzes operations and suggests improvements that can be made.

- **Business rules management**: Automating business rules and managing them based on inputs provided by subject matter experts.
- **Robotics**: Using software to imitate human behavior in the automation of actions and related interactions with software systems.

Decision management results in efficiency and productivity, two critical factors for successful business operations. As a concept, decision management can be used in a wide number of industries, functions, and areas of business. There are so many businesses that make scores of operational decisions on a daily basis. The quality of these decisions has a direct impact on the effectiveness of the company. All decisions are impacted by data, regulations, market dynamics, and decision management—and therefore becomes a necessity.

There are a range of ways that decision management interacts and influences a variety of sectors.

In the financial sector:

- Mortgage approval
- Insurance and loan disbursal
- Financial trading
- Leasing and billing

In the corporate world:

- Supply chain management
- Product configuration
- Risk management
- Cross-selling
- Logistics

In the public sector:

- Permit approval
- Subsidy determination
- Welfare
- Taxes

Benefits of Decision Management

For businesses considering decision management, there are a number of benefits.

Better Utilization of Time

Regardless of the model of the decision management support system, research shows that it reduces the decision time cycle. Employee productivity is the immediate benefit from the time saved.

Better Efficacy

The effectiveness of decisions made with decision management is still debated because the quality of these decisions is hard to measure. Research has largely taken up the approach of examining soft measures like a perceived decision quality instead of objective measures. Those who advocate the creation of data warehouses are of the strong opinion that better and larger scale analyses can definitely enhance decision-making.

Better Interpersonal Communication

Decision management systems open the door for better communication and collaboration among all decision makers. Set rules ensure that all decision makers are on a single platform, sharing facts and any assumptions made. Data-driven rule sets analyze and provide decision-makers with the best version of the possible outcome, encouraging fact-based decision-making. Better access to data always enhances the quality and clarity of decisions.

Cost Reduction

An outcome from good decision management rule sets is saving costs in labor (which comes from good decision-making, lowered infrastructure, and technological costs).

Better Learnings

In the long term, a by-product of decision management is that it encourages learning. There is more openness to new concepts, fact-based understanding of businesses, and the overall decision-making environment. Decision management can also come in handy to train new employees—an advantage yet explored in full.

Increased Organizational Control

With decision-making rule sets, a lot of transactional data is made available for constant performance checks and ad hoc enquiries by business heads. This gives management a better look at how business operations work. Managers find this to be a useful aspect of decision-making. There is a financial benefit to highly-detailed data, and this gradually becomes evident.

Disadvantages of Decision Management

As with any system, decision management systems can have a few disadvantages.

Information Overload

Considering the amount of data that goes through the system (and the fact that a problem is analyzed from multiple aspects), there are chances of information overload. With too many variables available on hand, the decision maker may be faced with a dilemma. Streamlined rule sets can help.

Over-Dependence

When decision making is completely computer based, it can lead to over-dependence. While it does free up man hours for better use of skills, it also increases dependency on computer-based decision making. Individuals can be less inclined to think independently and come to rely on computers to think for them.

Subjectivity

One of the important aspects of decision making is the number of alternatives that are offered based on objectivity. Subjectivity then tends to take a backseat, and this can affect decision-making and impact businesses. Things that cannot be measured cannot be factored.

Overemphasis on Decision Making

Not all issues an organization is faced with needs the power of decision management. An emphasis has to be placed on utilizing decision making capabilities for relevant issues.

Types of Decision Support Systems for Decision Making

Decision support systems are classified into two types:

1. **Model-Based Decision Support Systems**: These stand independent of any corporate information system. They work on the basis of strong theory or models and come with an excellent interface for easy interactivity.
2. **Data-Based Decision Support Systems**: These set-ups collect large amounts of data from a variety of sources, store it in warehouses, and analyze it. The warehouse stores historical data and also comes with some reporting and query tools.

In data-based decision support systems there are two main techniques that are employed:

1. **Online Analytical Processing (OLAP)**: Based on queries, this provides quick answers to some complex business needs. Managers and analysts can actively interact and examine data from multiple viewpoints.
2. **Data Mining**: By finding patterns and rules in existing data, useful decision making information can be extracted to help in trend and consumer behavior patterns.

There is no doubt that the most significant benefit of decision management is the overall improved management of revenue and profitability thanks to better decisions. While the investment in the system may be higher, the return on investment is undoubtedly quick. Taking away guesswork when it comes to decision-making ensures that it is a sensible business solution.

FOUNDATIONS AND TRENDS IN MACHINE LEARNING

NALLANNAGARI ANUSHA (22MCA26)
SABA TUNSEEM (22MCA33)

Foundations and Trends® in Machine Learning publishes exclusively long (± 100 pages) review and tutorial papers. Original research papers will be rejected.

Initial submission

If you intend to write, or are in the process of writing, a paper that fits within the format and scope of Foundations and Trends® in Machine Learning, we will be pleased to hear from you. In the first instance, send an abstract and table of contents for initial review to the publisher (see the journal homepage for contact details). After this initial submission, a preliminary acceptance may follow. The full draft paper will be subject to a reviewing process to ensure quality standards and balance before being finally accepted.

Manuscript submission

Once you have received preliminary acceptance, you can submit your full manuscript. Below are the manuscript preparation and submission instructions.

To provide rapid publication, we ask our authors to create their articles in LaTeX or Microsoft Word.We encourage you to use our FnT LaTeX style files but we also accept manuscripts in Word. The final output will be three versions of your article: a journal article, a book, and an e-book version. To help you further, please adopt the following guidelines when preparing your manuscript:

Abstract

Provide an abstract to the article of approximately 200 words. It should describe what the paper reviews and for whom it is of interest. The cover text in the printed book version will be based on the abstract. The abstract will also appear in various online and printed abstract journals.

References

A list of references must be provided at the end of the paper. The reference style for FnT Machine Learning is alphabetical. Please also ensure that every reference is cited in your text.

- It is very important that you supply as complete a reference as possible and that it is structured in the manner requested. Please see an example of the reference and citation style below:

Alphabetical reference style:

- Example:
  Reis, R. and P. C. Stocken (2007), "Strategic consequences of historical cost and fair value measurements". *Contemporary Accounting Research* 24(2), 557-584.
- Possible in-text citations for this reference are: "Reis and Stocken [2007]" or "[Reis and Stocken, 2007]".
- When there are three or more authors, the in-text citation will use "et al.". So, for example, the in-text citation for the following reference:
- Nickell, S., L. Nunziata, W. Ochel, and G. Quintini (2002), "The Beveridge-curve, unemployment and wages in the OECD from the 1960's to the 1990's". In: P. Aghion, R. Frydman, J. Stiglitz, and M. Woodford (eds.): *Knowledge, Information, and Expectation in Modern Macroeconomics*. Princeton: Princeton University Press, pp. 394-431.
- would be: "Nickell et al. [2002]" or "[Nickell et al., 2002]".

Reinforcement Learning: A Paradigm Shift in Artificial Intelligence

Akshitha CD(22MCA02)

Bhumika CR(22MCA06)

Abstract: Artificial Intelligence (AI) has made significant strides in mimicking human cognitive abilities, and one of the most promising areas within AI is reinforcement learning (RL). Reinforcement learning is a machine learning paradigm that enables agents to learn optimal actions by interacting with an environment. This article explores the fundamental concepts of reinforcement learning, its applications across various domains, and the challenges it poses. The synthesis of RL with real-world problems has led to ground breaking achievements, yet several obstacles remain, hindering its full potential. This article discusses these challenges and presents potential future directions for reinforcement learning research.

Introduction Reinforcement learning has revolutionized the way machines learn from experience, moving beyond traditional supervised and unsupervised learning techniques. The central idea of reinforcement learning is to train agents to make sequential decisions by providing feedback in the form of rewards or penalties. This article delves into the core components of reinforcement learning, including the agent, environment, actions, rewards, and policies. It explores how RL algorithms like Q-learning, policy gradients, and deep reinforcement learning have driven AI to achieve remarkable feats in domains such as robotics, gaming, and autonomous systems.

Key Concepts and Components: This section provides an in-depth analysis of the building blocks of reinforcement learning. It discusses Markov decision processes (MDPs) as the formal framework for modeling sequential decision-making problems, and explores how agents use exploration and exploitation strategies to learn optimal policies. The concept of value functions and Bellman equations are explained as crucial tools for estimating the long-term rewards associated with different actions. Additionally, the utilization of neural networks to approximate value functions in deep reinforcement learning is elaborated upon.

Applications of Reinforcement Learning: The versatility of reinforcement learning is demonstrated through its applications in diverse domains. In robotics, RL has enabled robots to learn complex tasks, such as grasping objects and walking, through trial and error. In the field of healthcare, RL aids in personalizing treatment plans and optimizing resource allocation in hospitals. Gaming showcases some of the earliest successes of RL, with agents mastering games like Go, Chess, and Dota 2, often outperforming human players. Self-driving cars exemplify the potential of RL in autonomous systems, enabling vehicles to make real-time decisions in dynamic environments.

Challenges and Open Problems: While reinforcement learning has achieved impressive milestones, several challenges persist. The problem of exploration versus exploitation remains a key concern, as agents must balance between trying out new actions and exploiting known optimal actions. The issue of sample inefficiency poses a challenge in domains where gathering real-world data is expensive or time-consuming. The instability of training deep RL models and the ethical implications of RL in decision-making processes also warrant attention.

Future Directions: The future of reinforcement learning holds immense promise. The development of more efficient exploration strategies, transfer learning techniques, and methods to enhance sample efficiency could pave the way for wider adoption of RL in various domains. Adapting RL to multi-agent systems and addressing safety concerns in real-world applications are important avenues for future research. Additionally, exploring the synergy between reinforcement learning and other AI approaches like meta-learning and imitation learning could lead to novel breakthroughs.

Conclusion: Reinforcement learning stands as a paradigm shift in the field of artificial intelligence, enabling machines to learn optimal actions through interaction with their environments. The remarkable achievements in domains ranging from robotics to gaming highlight its potential. However, challenges such as exploration-exploitation trade-offs and sample inefficiency remain. As the field continues to evolve, addressing these challenges and exploring new directions will drive the integration of reinforcement learning into our daily lives, reshaping the way we interact with technology and furtheringthe boundaries of AI.

References:

1.Ren B, Ye X, Pan Z and Zhang T (2022). Versatile Control of Fluid-directed SolidObjects Using Multi-task Reinforcement Learning, ACM Transactions on Graphics, 2.Movin M, Junior G, Hollmén J and Papapetrou P. Explaining Black Box Reinforcement Learning Agents Through Counterfactual Policies. Advances in Intelligent Data AnalysisXXI.

3.Gao C, Huang K, Chen J, Zhang Y, Li B, Jiang P, Wang S, Zhang Z and He X. Alleviating Matthew Effect of Offline Reinforcement Learning in Interactive Recommendation. Proceedings of the 46th International ACM SIGIR Conference on Research and Development in Information Retrieval.

NATURAL LANGUAGE PROCESSING

Haritha V (22MCA13)
Divyashree D (22MCA09)

Abstract:-

The advances in the digital era have necessitated the adoption of communication as the main channel for modern business. In the past, business negotiations, profiling, seminars, shopping and agreements were in-person but today everything is almost digitalized**.**

Introduction:-

Natural Language Processing (NLP) is a vibrant and rapidly evolving field within the broader domain of the Artificial Intelligent (AI). It focuses on developing algorithms and models that enables computer to understand, generate and interact with Human language.

This document provides a concise overview of prominent NLP eJournals that contribute to the advancement of AI and NLP.

1. **Computational Linguistics:** As the official journal of the Association for Computational linguistics (ACL), Computational linguistics is a premier eJournal in the NLP and Ai Domain. It covers a wide spectrum of topics, including machine translation, sentiment analysis, language generation, and more.
2. **Transactions of the Association for Computational linguistics (TACL):-** TACL, open-access eJournal, publishes articles in various areas of Computational linguistics and NLP. It emphasizes both theoretical insights and practical applications, fostering a balanced view of the field's of developments.
3. **IEEE Transactions on Natural language Processing:-** This IEEE ejournal provides a platform for researchers to publish their cutting-edge work in NLP and AI. It offers a broad scope, encompassing topics such as semantic analysis, discourse modeling and dialogue Systems.
4. **Natural Language Engineering:-** Focused on the interaction between linguistic studies and computational techniques, natural language engineering bridges the gap between theoretical linguistics and NLP Applications. It is a valuable resource for researchers interested in practical aspects of NLP.
5. **Language Resources and Evaluation:-** Language resources and evaluation is a significant ejournal that emphasizes the importance of the linguistic resources, such as annotated datasets and corpora, in NLP research. It showcases methodologies for resource creation and evaluation. Contributing to the foundation of NLP advancements.
6. **Journal of Machine Learning Research (JMLR):-** While primarily centered on machine learning, JMLR often features research that intersects with NLP and AI. This eJournal is ideal for researchers exploring the integration of machine learning techniques into NLP tasks.
7. **Empirical methods in Natural Language Processing (EMNLP):-** EMNLP is a prominent conference in the NLP field, and its proceeding are often published as and eJournal source. It covers a broad spectrum of NLP topics, including information retrieval, text classification, and dialogue systems.

Conclusion:

In the dynamic landscape of AI and NLP, staying informed about the latest research is essential for professionals and researchers alike. These eJournals provide valuable insights into the advancements, challenges and trends within the NLP fields. By regularly engaging with these publication, individuals can contribute to and benefit from the ongoing progress in AI and NLP research.

GENERATIVE AI

Anjushree

Harishree

Introduction

In recent years, the field of artificial intelligence has witnessed remarkable advancements, with one of the most fascinating branches being Generative AI. Generative AI, also known as Generative Adversarial Networks (GANs) and Variational Autoencoders (VAEs), has revolutionized the way we think about creativity, art, and even problem-solving. This ejournal aims to delve into the captivating world of Generative AI, exploring its applications, challenges, and the profound impact it has on various industries.

Understanding Generative AI

Generative AI is a subset of artificial intelligence that focuses on the creation of data, such as images, text, music, and more. Unlike traditional AI systems that rely on pre-defined rules and data, Generative AI is capable of generating new content autonomously. It does this by learning patterns and structures from existing data andthen producing novel output based on those learned patterns.

Applications of Generative AI

Art and Creativity: Generative AI has been used extensively in the world of art. It cancreate stunning paintings, compose music, and even generate poetry. Artists and musicians are collaborating with AI systems to explore new dimensions of creativity.

Content Generation: In the field of content creation, Generative AI can assist in generating written content, including articles, reports, and even code. This can save time and effort for content creators and developers.

Healthcare: Generative AI is making significant strides in healthcare, aiding in the generation of 3D medical images, drug discovery, and even predicting patient outcomes based on electronic health records.

Design and Fashion: Generative AI can assist in the creation of fashion designs, product prototypes, and architectural blueprints. It enables designers to explore a vast array of possibilities quickly.

Gaming: Video game developers use Generative AI to create realistic environments, characters, and even generate in-game content, leading to more immersive gaming experiences.

Challenges and Ethical Considerations

While Generative AI offers tremendous potential, it also comes with challenges and ethical considerations. Some of the key challenges include:

Bias and Fairness: AI models can inherit biases from the data they are trained on, leading to biased outputs. Ensuring fairness and addressing biases is a significant challenge.

Quality Control: Not all generated content is of high quality. Ensuring that the AI produces valuable and accurate output remains a challenge.

Data Privacy: The use of Generative AI raises concerns about data privacy, especiallywhen generating content based on personal data.

Ethical Use: Determining the ethical boundaries of Generative AI, particularly in generating deepfakes and potentially harmful content, is an ongoing debate.

Conclusion

Generative AI represents a remarkable leap forward in the field of artificial intelligence, offering boundless possibilities for creativity and problem-solving. Its applications span various industries, from art and healthcare to content generation and design. However, it is essential to navigate the challenges and ethical considerations associated with this technology carefully. As Generative AI continues to evolve, it willundoubtedly reshape the way we think about creativity and the potential of artificial intelligence.

Artificial Intelligence (VISION)

Ashmika Shandilya
Javeriya R

Vision AI (also known as Computer Vision) is a field of computer science that trains computers to replicate the human vision system. This enables digital devices (like face detectors, QR Code Scanners) to identify and process objects in images and videos, just like humans do.

How does computer vision work?

Computer vision needs lots of data. It runs analyses of data over and over until it discerns distinctions and ultimately recognize images. For example, to train a computer to recognize automobile tires, it needs to be fed vast quantities of tire images and tire-related items to learn the differences and recognize a tire, especially one with no defects.

Applications of Computer Vision AI :

1. Image Segmentation

Its a process of partitioning an image from multiple regions and pieces, based on pixel characteristics in an image. Generally used for examining purposes, image segmentation involves separating foreground from background or clustering parts of an image by pixels, based on similarity in color or shape. The image shown below exemplifies image segmentation, where parts of the image are differentiated with colors.

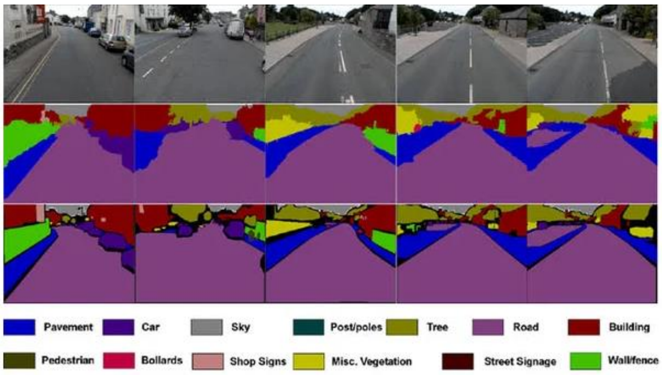

2. Object Detection

This field of computer vision AI deals with the detection of one or multiple objects in an image or a video. For example, surveillance cameras smartly detect humans and their activities (no movement, objects like guns or knife, etc.) so that caution is passed for suspicious activities.

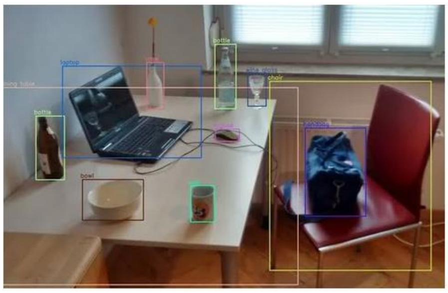

3. Facial Recognition

The facial recognition technique aims at detecting an object or human face in the image. It is one of the complex applications of computer vision because of variability in human faces-expression, pose, skin color, the difference in camera quality, position or orientation, image resolution, etc. However, this technique is prominently used. Smartphones use it for user authentication. Facebook uses the same technique when it gives tagging suggestions for people in a picture

4. Edge Detection

Edge detection deals with finding the boundaries of objects within an image. This is done by detecting discontinuities in brightness. Edge detection can be a great help in data extraction and image segmentation.

NEUROMORPHIC COMPUTING

Vandana and Renuka

Neuromorphic computing is a method of computer engineering in which elements of a computer are modeled after systems in the human brain and nervous system. The term refers to the design of both hardware and software computing elements. Neuromorphic computing is sometimes referred to as neuromorphic engineering.

Neuromorphic engineers draw from several disciplines -- including computer science, biology, mathematics, electronic engineering and physics -- to create bio-inspired computer systems and hardware. Of the brain's biological structures, neuromorphic architectures are most often modelled after neurons and synapses. This is because neuroscientists consider neurons the fundamental units of the brain.

Neuromorphic computers have the following characteristics:

- **Collocated processing and memory.** The brain-inspired neuromorphic computer chips process and store data together on each individual neuron instead of having separate areas for each. By collocating processing and memory, neural net processors and other neuromorphic processors avoid the von Neumann bottleneck and can have both high performance and low energy consumption at the same time.
- **Massively parallel.** Neuromorphic chips, such as Intel Lab's Loihi 2, can have up to one million neurons. Each neuron operates on different functions simultaneously. In theory, this lets neuromorphic computers perform as many functions at one time as there are neurons. This type of parallel functioning mimics stochastic noise, which is the seemingly random firings of neurons in the brain. Neuromorphic computers are designed to process this stochastic noise better than traditional computers.
- **Inherently scalable.** Neuromorphic computers do not have traditional roadblocks to scalability. To run larger networks, users add more neuromorphic chips, which increases the number of active neurons.
- **Event-driven computation.** Individual neurons and synapses compute in response to spikes from other neurons. This means only the small portion of neurons actually

processing spikes are using energy; the rest of the computer remains idle. This makes for extremely efficient use of power.

- **High in adaptability and plasticity.** Like humans, neuromorphic computers are designed to be flexible to changing stimuli from the outside world. In the spiking neural network -- or SNN -- architecture, each synapse is assigned a voltage output and adjusts this output based on its task. SNNs are designed evolve different connections in response to potential synaptic delays and a neuron's voltage threshold. With increased plasticity, researchers hope neuromorphic computers will learn, solve novel problems and adapt to new environments quickly.
- **Fault tolerance.** Neuromorphic computers are highly fault tolerant. Like the human brain, information is held in multiple places, meaning the failure of one component does not prevent the computer from functioning.

Challenges of neuromorphic computing

Many experts believe neuromorphic computing has the potential to revolutionize the algorithmic power, efficiency and capabilities of AI as well as reveal insights into cognition. However, neuromorphic computing is still in early stages of development, and it faces several challenges:

- **Accuracy.** Neuromorphic computers are more energy efficient than deep learning and machine learning neural hardware and edge graphics processing units (GPUs). However, they have still not proven themselves to be conclusively more accurate than them. Combined with the high costs and complexity of the technology, the accuracy issue leads many to prefer traditional software.
- **Limited software and algorithms.** Neuromorphic computing software has not caught up with the hardware. Most neuromorphic research is still conducted with standard deep learning software and algorithms developed for von Neumann hardware. This limits research results to standard approaches, which neuromorphic computing is trying to evolve beyond. Katie Schuman, a neuromorphic computing researcher and an assistant professor at the University of Tennessee, said in an interview with *Ubiquity* that adoption of neuromorphic computing technologies "will require a paradigm shift in how we think about computing as a whole …. Though this is a difficult task, continued innovation in computing depends on our willingness to move beyond our traditional von Neumann systems."

- **Inaccessible.** Neuromorphic computers aren't available to nonexperts. Software developers have not yet created application programming interfaces, programming models and languages to make neuromorphic computers more widely available.
- **Benchmarks.** Neuromorphic research lacks clearly defined benchmarks for performance and common challenge problems. Without these standards, it's difficult to assess the performance of neuromorphic computers and prove efficacy.
- **Neuroscience.** Neuromorphic computers are limited to the known structures of human cognition, which is still far from completely understood. For instance, there are several theories that propose human cognition is based on quantum computation, such as the Orch (OR) theory proposed by Sir Roger Penrose and Stuart Hameroff. If cognition requires quantum computation as opposed to standard computation, neuromorphic computers would be incomplete approximations of the human brain and might need to incorporate technologies from fields like probabilistic and quantum computing.

Neuromorphic technology is expected to be used in the following ways:

- deep learning applications;
- next-generation semiconductors;
- transistors;
- accelerators; and
- autonomous systems, such as robotics, drones, self-driving cars and artificial intelligence (AI).

Internet of Things

Yeseera Farhath S
Yashashvini R

Internet of Things is referred to as Iot. It describes the connectivity of physical objects that are implanted with connectivity, software, and sensors that allow them to connect and exchange data. Examples of such products are appliances and cars.

This technology makes it possible to gather and share data from a huge network of gadgets, opening the door to the development of automatic and more effective systems.

The Internet of Things (iot) is the networking of physical items with electronics built into its architecture to enable communication and the detection of interactions between them or with the surrounding environment.

Iot-based technology will provide higher levels of services in the future years, effectively altering how individuals go about their daily lives. Just a few categories where iot is well established include improvements in medicine, power, gene therapies, agriculture, smart cities, and smart homes.

The Internet of Things (iot) is a network of networked computing devices that are implanted in commonplace things and allow them to send and receive data.

There are two approaches to developing iot:

1. Create a separate online project with only tangible items.
2. Expand the Internet further, yet doing so calls for advanced technology like quick huge data storage and strict cloud computing, both of which are pricy.

Utilizing iot Devices:

Data collection and transmission: Sensors are frequently used for this purpose and are utilized based on application-specific needs.

Depending on signals generated by sensors or processing units, activate the device: Actuator devices indicate what action should be taken when specific conditions are met or when a trigger is engaged in accordance with user requirements.

Receive Information: Users or devices can obtain specific information from network devices for analysis and processing needs.

Communication Support: The phenomena of communication between two networks or between two or more iot devices of the same or different networks is known as communication assistance.

Different communication protocols, such as MQTT, Constrained Application Protocol, zigbee, FTP, HTTP, etc., can be used to do this.

Modern Applications:

- Smart Grids and energy saving
- Smart cities
- Smart homes/Home automation
- Healthcare
- Earthquake detection
- Radiation detection/hazardous gas detection
- Smartphone detection

Iot Advantages:

- Increased productivity and task automation.
- Increased information convenience and availability.
- Improved system and device monitoring and control.
- Increased capacity for data collection and analysis.
- More effective decision-making.
- Cost reductions.

Iot Disadvantages:

- Concerns about security and the possibility of hacking or data breaches.
- Concerns about privacy relating to the use and collection of personal information.
- Technology dependence and the risk of system failure.
- Device interoperability and standardization are limited.
- Complexity and more frequent maintenance demands.
- High upfront investment expenses.
- Some devices only have a short battery life.

DEEP LEARNING

Divya N M

Likitha

Deep learning drives many artificial intelligence (AI) applications and services that improve automation, performing analytical and physical tasks without human intervention. Deep learning technology lies behind everyday products and services (such as digital assistants, voice-enabled TV remotes, and credit card fraud detection) as well as emerging technologies (such as self-driving cars).

How Deep Learning works?

Deep learning uses neural networks with multiple layers to automatically learn patterns from data. It starts with random weights, passes data through the network, computes a loss (error), adjusts weights to minimize the loss using backpropagation, and iterates this process to improve predictions. Finally, the trained model can make accurate predictions on new, unseen data.

1. Data Input: Deep learning starts with a dataset containing input data (e.g., images, text, or numerical values) and corresponding target outputs (e.g., labels, categories, or values).

2. Neural Network Architecture: A neural network is designed with layers of interconnected nodes or neurons. Typically, there are input, hidden, and output layers. The depth of the network (number of hidden layers) distinguishes deep learning.

3. Weight Initialization: The network's connections (synaptic weights) are initialized with random values. These weights determine the strength of connections between neurons.

4. Forward Propagation: During training or inference, data flows forward through the network. Each neuron receives input from the previous layer, computes a weighted sum, addsa bias, and passes the result through an activation function.

5. Activation Functions: Activation functions introduce non-linearity, allowing the network tomodel complex relationships in data.

6. Loss Function: A loss function measures the difference between the network's predictions and the actual target values. It quantifies the model's performance.

7. Backpropagation: Backpropagation calculates gradients of the loss with respect to the model's parameters (weights and biases) in reverse order, starting from the output layer. These gradients guide weight updates to minimize the loss.

8. Training Iteration: The training process involves iteratively passing data through the network, calculating the loss, and adjusting weights using optimization algorithms (e.g., gradient descent). This process continues until the model's performance converges.

9. Validation and Testing: After training, the model is evaluated using a separate validation dataset to ensure it generalizes well. Testing is performed on an unseen test dataset to assess real-world performance.

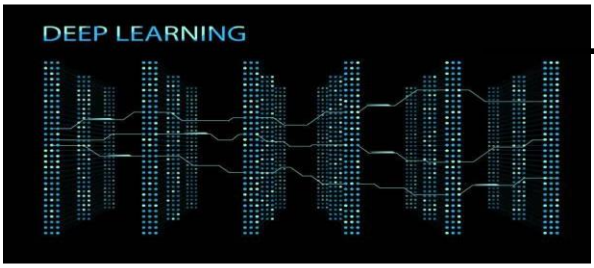

Application:

1. <u>Gaming:</u> Deep reinforcement learning is used in gaming AI to develop agents that can learn to play games like chess, Go, and video games.

2. <u>Healthcare:</u> Deep learning is applied in medical image analysis for tasks like detecting tumours' in medical scans (MRI, CT), classifying skin lesions, and predicting disease progression.

3. <u>Manufacturing and Industry:</u> Deep learning helps optimize manufacturing processes, predict equipment failures, and improve quality control in industries like automotive and electronics.

4. <u>Robotics:</u> Deep learning is crucial for robot perception, control, and manipulation tasks, making robots more autonomous and versatile.

Conclusion

Deep learning has redefined AI, achieving remarkable success in various domains. Its data-hungry nature demands large datasets, and complex models can be computationally intensive. Ethical concerns about bias and privacy require careful consideration. Transfer learning accelerates progress by leveraging pertained models. Despite challenges, deep learning's interdisciplinary impact continues to shape our technological landscape.